



(19)

Europäisches Patentamt
Europ an Patent Office
Office européen des brevets



(11) EP 0 458 859 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:

30.07.1997 Bulletin 1997/31

(51) Int Cl. 6: G10L 5/04

(21) Application number: 90903452.2

(86) International application number:
PCT/US90/00528

(22) Date of filing: 02.02.1990

(87) International publication number:
WO 90/09657 (23.08.1990 Gazette 1990/20)(54) TEXT TO SPEECH SYNTHESIS SYSTEM AND METHOD USING CONTEXT DEPENDENT
VOWELL ALLOPHONESSYSTEM UND METHODE ZUR TEXT-SPRACHE-UMSETZUNG MIT HILFE VON
KONTEXTABHÄNGIGEN VOKALALLOPHONENSYSTEME ET PROCEDE DE SYNTHESE DE TEXTE EN PAROLES UTILISANT DES
ALLOPHONES DE VOYELLE DEPENDANT DU CONTEXTE(84) Designated Contracting States:
DE GB(56) References cited:
US-A- 4 627 001 US-A- 4 831 654

(30) Priority: 17.02.1989 US 312692

- PROCEEDINGS OF THE NATIONAL COMMUNICATIONS FORUM. vol. 43, no. 2, 2 October 1989, OAK BROOK, ILLINOIS US pages 1104 - 1108; B. MALSHEEN: 'Recent developments in text to speech research'
- INTERNATIONAL CONFERENCE ON ACOUSTICS SPEECH AND SIGNAL PROCESSING 2 April 1979, WASHINGTON DC, USA pages 891 - 894; SCHWARTZ ET AL: 'Diphone synthesis for phonetic vocoding'
- INTERNATIONAL CONFERENCE ON ACOUSTICS SPEECH AND SIGNAL PROCESSING vol. 1, 11 April 1988, NEW YORK, USA pages 659 - 662; NAKAJIMA, HAMADA: 'Automatic generation of synthesis units based on context oriented clustering'
- IEEE ACOUSTICS, SPEECH, AND SIGNAL PROCESSING MAGAZINE. 1 April 1984, NEW YORK US pages 4 - 29; GRAY: 'Vector quantization'

(43) Date of publication of application:
04.12.1991 Bulletin 1991/49(73) Proprietor: CENTIGRAM COMMUNICATIONS
CORPORATION
San Jose, California 95139 (US)(72) Inventors:

- MALSHEEN, Bathsheba, J.
San Francisco, CA 94117 (US)
- GRONER, Gabriel, F.
Palo Alto, CA 94306 (US)
- WILLIAMS, Linda, D.
San Jose, CA 95117 (US)

(74) Representative: Cross, Rupert Edward Blount
BOULT WADE TENNANT,
27 Fenchurch Street
London EC4A 1PQ (GB)

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

The present invention relates generally to speech synthesis, and particularly to methods and systems for converting textual data into synthetic speech.

5

BACKGROUND OF THE INVENTION

10

The automatic conversion of text to synthetic speech is commonly known as text to speech (TTS) conversion or text to speech (TTS) synthesis. A number of different techniques have been developed to make TTS conversion practical on a commercial basis. An excellent article on the history of TTS development, as well as the state of the art in 1987, is Dennis H. Klatt, *Review of text-to-speech conversion for English*, *Journal of the Acoustical Society of America* vol. 82(3), September 1987.

15

A number of commercial products use TTS techniques, including the Speech Plus Prose 2000 (made by the assignee of the applicants), the Digital Equipment DECTalk, and the Infovox SA-101.

Overview of Prior Art TTS

20

Referring to Figure 1 most commercial TTS products first convert text into a stream of phonemes (with representations for emphasis and stress) and then use a "synthesis by rule" technique for converting the phonemes into synthetic speech. For example, in the Speech Plus Prose 2000 Text-to-Speech Converter the first step of the TTS process is text normalization (box 20), which expands abbreviations to their full word form. The Text Normalization routine 20 expands numbers, monetary amounts, punctuation and other non-alphabetic characters into their full word equivalents.

25

Most words are converted to phonemes by a set of Word to Phoneme Rules 24. However, the pronunciation of some words do not follow the standard rules. The phoneme strings for these special words are stored in a Dictionary Look-Up Table 22. In a typical TTS system, 3000 to 5000 such words are stored in the Dictionary 22. Thus, using either the Dictionary 22 or the Phoneme Rules 24 for each particular word, all text input is converted into phoneme strings.

The Word-Level Stress Assignment routine 26 assigns stress to phonemes in the phoneme string. Variations in assigned stress result in pitch and duration differences that make some sounds stand out from others.

30

It is well known that the pronunciation of phonemes in human (or natural) speech is context dependent. To mimic natural speech, the synthetic pronunciation of each phoneme is determined by a set of rules which analyze the phonetic context of the phoneme. The Allophones routine 28 assigns allophones to at least a portion of the consonant phonemes in the phoneme string 25.

35

Allophones are variants of phonemes based on surrounding speech sounds. For instance, the aspirated "p" of the word *pit* and the unaspirated "p" of the word *spit* are both allophones of the phoneme "p".

One way to try to make synthetic speech sound more natural is to "assign" or generate allophones for each phoneme based on the surrounding sounds, as well as the speech rate, syntactic structure and stress pattern of the sentence. Some prior art TTS products, such as the Speech Plus Prose 2000, assign allophones to certain consonant phonemes based on the context of those phonemes. In other words, an allophone is selected for a particular consonant phoneme based on the context of that phoneme in a particular word or sentence.

40

The Sentence-Level Prosodics rules 30 in the Speech Plus Prose 2000 determine the duration and fundamental frequency pattern of the words to be spoken. The resultant intonation contour gives sentences a semblance of the rhythm and melody of a human speaker. The prosodics rules 30 are sensitive to the phonetic form and the part of speech of the words in a sentence, as well as the speech rate and the type of the prosody selected by the user of the system.

45

The Parameter Generator 40 accepts the phonemes specified by the early portions of the TTS system, and produces a set of time varying speech parameters using a "constructive synthesis" algorithm. In other words, an algorithm is used to generate context dependent speech parameters instead of using pieces of prestored speech. The purpose of the constructive synthesis algorithm is to model the human vocal tract and to generate human sounding speech.

50

The speech parameters generated by the Parameter Generator 40 control a digital signal processor known as a Formant Synthesizer 42 because it generates signals which mimic the formants (i.e., resonant frequencies of the vocal tract) characteristic of human speech. The Formant Synthesizer outputs a speech waveform 44 in the form of an electrical signal that is used to drive a audio speaker and thereby generates audible synthesized speech.

55

Diphone Concatenation. Another technique for TTS conversion is known as diphone concatenation. A diphone is the acoustic unit which spans from the middle of one phoneme to the middle of the next phoneme. TTS conversion systems using diphone concatenation employ anywhere from 1000 to 8000 distinct diphones. In diphone concatenation systems, each diphone is stored as a chunk of encoded real speech recorded from a particular person. Synthetic speech is generated by concatenating an appropriate string of diphones. Due to the fact that each diphone is a fixed package of encoded real speech, diphone concatenation has difficulty synthesizing syllables with differing stress and

timing requirements. While some experimental diphone concatenation systems have good voice qualities, the inherent timing and stress limitations of concatenation systems have limited their commercial appeal. Some of the limitations of diphone concatenation systems may be overcome by increasing the number of diphones used so as to include similar diphones with different durations and fundamental frequencies, but the amount of memory storage required 5 may be prohibitive.

A similar technique, called demisyllable concatenation employs demisyllables instead of diphones. A demisyllable is the acoustic unit which spans from the start of a consonant to the middle of the following vowel in a syllable, or from the middle of a vowel to the end of the following consonant in a syllable.

One reason for the prevalence of TTS systems which use "synthesis by rule" techniques, as opposed to diphone 10 or demisyllable concatenation systems, is that synthesis by rule provides a greater ability to vary timing, intonation and allophonic detail - all of which are important to making synthetic speech intelligible, variable and pleasant to listen to. In addition, it has been demonstrated that the synthesis of phonemes follows certain patterns that can be generalized and represented by a set of rules.

Generally, diphone concatenation systems and synthesis by rule systems have different strong points and weaknesses. Diphone concatenation systems can sound like a person when the proper diphones are used because the 15 speech produced is "real" encoded speech recorded from the person that the system is intended to mimic. Synthesis by rule systems are more flexible in terms of stress, timing and intonation, but have a machine-like quality because the speech sounds are synthetic.

An embodiment of the present invention can be thought of as a hybrid of the synthesis by rule and diphone concatenation techniques. 20

Instead of using encoded (i.e., stored real speech) diphones, the present invention incorporates into a synthesis by rule system vowel allophones that are synthetic, but which resemble the full allophonic repertoire of a particular person.

25 Vowel Allophones

To a large degree, the prior art TTS systems and techniques generate allophones only for consonant phonemes. Vowel phonemes are generally given a static representation (i.e., are represented by a fixed set of formant frequency and bandwidth values), with "allophones" being formed by "smoothing" the vowel's formants with those of the neighboring phonemes. 30

More precisely, the fixed representation of each vowel phoneme is a partial set of formant frequency and bandwidth values which are derived by analyzing and selecting or averaging the formant values of one or more persons when speaking words which include that vowel phoneme. Vowel allophones (i.e., context dependent variations of vowel phonemes) are generated in the prior art systems, if they are generated at all, by formant smoothing. Formant smoothing 35 is a curve fitting process by which the back and forward boundaries of the vowel phoneme (i.e., the boundaries between the vowel phoneme and the prior and following phonemes) are modified so as to smoothly connect the vowel's formants with those of its neighbors.

One embodiment of the present invention, on the other hand, stores an encoded form of every possible allophone in the English (or any other) language. While this would appear to be impractical, at least from a commercial viewpoint, 40 the embodiment provides a practical method of storing and retrieving every possible vowel allophone. More specifically, a vowel allophone library is used to store distinct allophones for every possible vowel context. When synthesizing speech, each vowel phoneme is assigned an allophone by determining the surrounding phonemes and selecting the corresponding allophone from the vowel allophone library.

The inventors have found that using a large library of encoded vowel allophones, rather than a small set of static 45 vowel phonemes, greatly improves the intelligibility and naturalness of synthetic speech. It has been found that the use of encoded vowel allophones reduces the machine-like quality of the synthetic speech generated by TTS conversion.

In the context of Figure 1, the inventors have improved the parameter generator 40 of the prior art Speech Plus Prose 50 2000 system by adding a vowel allophone capability. Thus the generation of vowel allophones is handled separately from the generation of consonant allophones by Allophonics module 28.

More generally, though, the invention does not depend on the exact TTS technique being used in that it provides a system and method for replacing the static vowel phonemes in prior art TTS systems with context dependent vowel allophones.

It is therefore a primary object of one embodiment of the present invention to improve the quality and intelligibility 55 of the synthetic speech produced by TTS conversion systems.

An object of another embodiment of the present invention is to improve the quality and intelligibility of synthetic speech produced by TTS conversion systems by generating context dependent vowel allophones.

An object of another embodiment of the present invention is to provide a large library of vowel allophones and a

technique for assigning allophones in the library to the vowel phonemes in a phrase that is to be synthetically enunciated, so as to generate natural sounding vowel phonemes.

An object of another embodiment of the present invention is to provide a TTS conversion system that sounds like a particular person. A related object is provide a methodology for adapting TTS conversion systems to make them sound like particular individuals.

An object of yet another embodiment of the present invention is to provide a practical method and system for storing and retrieving a large library of vowel allophones, representing all or practically all of the vowel allophones in a particular language, so as enable use of the present invention in commercial applications.

10 SUMMARY OF THE INVENTION

The present invention provides a text-to-speech synthesis system, comprising text conversion means for converting a specified text into a corresponding string of consonant and vowel phonemes, each said phoneme being selected from a predefined set of phonemes including a multiplicity of consonant phonemes and a multiplicity of vowel phonemes; parameter generating means for generating speech parameters corresponding to said string of phonemes; and speech synthesizing means for generating a speech waveform corresponding to the speech parameters generated by said parameter generating means; characterized by vowel allophone storage means storing a multiplicity of predefined vowel allophones, each vowel allophone being represented by a set of speech parameters; said vowel allophones including allophones for a multiplicity of vowel phonemes; vowel phoneme to allophone conversion means, coupled to said text conversion means and said vowel allophone storage means, for computing a phoneme context value for each of at least a subset of said vowel phonemes in said string of phonemes, said phoneme context value comprising a function of the phonemes in said string of phonemes which precede and follow said vowel phoneme, and for then assigning to said vowel phoneme a selected one of said predefined vowel allophones corresponding to said computer phoneme context value; said parameter generating means including means for generating speech parameters for said assigned vowel allophones.

The present invention also provides a method of converting text strings into synthetic speech, the steps comprising: defining a set of phonemes, including a multiplicity of consonant phonemes and a multiplicity of vowel phonemes; converting a specified text string into a corresponding string of phonemes, said string of phonemes including consonant and vowel phonemes, each said phoneme being selected from said defined set of phonemes; and converting said string of phonemes into speech parameters and then generating an audio waveform corresponding to said speech parameters; characterized by: storing a multiplicity of predefined vowel allophones, each vowel allophone being represented by a set of speech parameters; for each of at least a subset of said vowel phonemes in said string of phonemes, computing a phoneme context value for said vowel phoneme as a function of the phonemes in said string of phonemes which precede and follow said vowel phoneme, and then assigning to said vowel phoneme a selected one of said predefined vowel allophones corresponding to said computed phoneme context value; and said converting step including converting said assigned vowel allophones into speech parameters which are then used to generate an audio waveform corresponding to said speech parameters.

40 BRIEF DESCRIPTION OF THE DRAWINGS

40 Embodiments of the present invention will now be described with reference to the drawings, in which:-

Figure 1 is a flow chart of the text to speech conversion process.

Figure 2 is a block diagram of a system for performing text to speech conversion.

Figure 3 depicts a spectrogram showing one vowel allophone.

Figure 4 depicts one formant of a vowel allophone.

Figure 5 is a block diagram of one formant code book and an allophone with a pointer to an item in the code book.

Figure 6 is a block diagram of the vector quantization process for generating a code book of vowel allophone formant parameters.

Figure 7A, 7B and 7C are block diagrams of the process for generating the formant parameters for a specified vowel allophone.

Figure 8 is a block diagram of an allophone context map data structure and a related duplicate context map.

Figure 9 is a block diagram of a vowel context data table.

Figure 10 is a block diagram of an alternate LLRR vowel context table.

Figure 11 is a block diagram of the process for generating speech parameters for a specified vowel allophone in an alternate embodiment of the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENT

Referring to Figure 2, the preferred embodiment of the present invention is a reprogrammed version of the Speech Plus Prose 2000 product, which is a TTS conversion system 50. The basic components of this system are a CPU controller 52 which executes the software stored in a Program ROM 54. Random Access Memory (RAM) 56 provides workspace for the tasks run by the CPU 52. Information, such as text strings, is sent to the TTS conversion system 50 via a Bus Interface and I/O Port 58. These basic components of the system 50 communicate with one another via a system bus 60, as in any microcomputer based system.

Note that boxes 20 through 40 in Figure 1 comprise a computer (represented by boxes 52, 54 and 56 in Figure 2) programmed with appropriate TTS software. It is also noted that the TTS software may be downloaded from a disk or host computer, rather than being stored in a Program ROM 54.

Also coupled to the system bus 60 is a Formant Synthesizer 62, which is a digital signal processor that translates formant and other speech parameters into speech waveform signals that mimic human speech. The digital output of the Formant Synthesizer 62 is converted into an analog signal by a digital to analog converter 64, which is then filtered by a low pass filter 66 and amplified by an audio amplifier 68. The resulting synthetic speech waveform is suitable for driving a standard audio speaker.

One embodiment of the present invention synthesizes speech from text using a variation of the process shown in Figure 1. In the preferred embodiment vowel allophones are assigned to vowel phonemes by an improved version of the parameter generator 40. In terms of the sequence of process steps, the vowel allophone assignment process takes place between blocks 30 and 40 in Figure 1.

As explained above, one embodiment of the present invention generates improved synthetic speech by replacing the fixed formant parameters for vowel phonemes used in the prior art with selected formant parameters for vowel allophones. The vowel allophones are selected on the basis of the "context" of the corresponding phoneme - i.e., the phonemes preceding and following the vowel phoneme that is being processed.

To understand the magnitude of this task, consider the following. Assume for the purposes of this example that the context of a vowel phoneme is defined solely by the phonemes immediately preceding and following the vowel phoneme. The preferred embodiment of the invention uses 57 phonemes (including 23 vowel phonemes, 33 consonant phonemes, and silence). For each vowel (i.e., vowel phoneme) there are 3136 (i.e., 56×56) possible phoneme-vowel-phoneme (PUP) contexts. In other words, there are 3136 possible allophones for each of the 23 vowel phonemes, or a total of 72,128 vowel allophones.

In the preferred embodiment, and many commercial products, the enunciation of a vowel phoneme is represented by four formants, requiring approximately 40 bytes to store each vowel allophone. The data structure for storing a single phoneme enunciation (i.e., allophone) is described in more detail below. Without using some form of data compression, it would require nearly three megabytes of memory to store the 72,128 possible vowel allophones. In most commercial applications, it is currently not practical to use so much memory just to store a library of vowel allophones. It should be noted that in many commercial applications, a TTS system is an "add-on board" which must occupy a relatively small amount of space and must cost less than a typical desktop computer.

One embodiment of the present invention provides a practical and relatively low cost method of storing and accessing the data for all 72,128 vowel allophones, using allophone data tables which occupy about one tenth of the space which would be required in a system that did not use data compression. Before explaining how this is done, it is first necessary to review the data used to represent vowel allophones.

Speech Formant Parameters

Figure 3 shows a somewhat simplified model of the speech spectrogram 80 for one vowel allophone. The speech spectrogram 80 shows four formants f_1 , f_2 , f_3 and f_4 . As shown, each formant has a distinct frequency "trajectory", and a distinct bandwidth which varies over the duration of the allophone. The frequency trajectory and bandwidth of each formant directly correlate with the way that formant sounds.

To store and retrieve any sound, one can simply record the soundwave and play it back. However, that is not practical when building a library of over 72,000 allophones because of the huge volume of memory which would be required to store the digital samples.

Rather, speech waveforms can be reconstructed from information stored in a much more compressed form because of knowledge about their structure and production. In particular, one standard method of reconstructing a speech waveform is to record the frequency trajectory of each formant, plus the bandwidth trajectory of at least the lower two or three formants. Then the waveform is synthesized by using the frequency and bandwidth trajectories to control a formant synthesizer. This method works because the formant frequencies are the resonant frequencies of the vocal tract and they characterize the shape of the vocal tract as it changes to produce the speech waveform.

Referring to Figures 3 and 4, in the present invention each individual allophone formant is represented by six

EP 0 458 859 B1

frequency measurements (bbx, v1x, v2x, v3x, v4x and fbx), four time measurements (t1x, t2x, t3x and t4x), and three bandwidth measurements (b3x, b5x and b7x), where "x" identifies the formant. These measurements trace the frequency trajectory of the formant, as well as changes in its bandwidth.

5 Table 1 lists the measurement parameters for a single allophone formant and describes the measured quantity represented by each parameter.

Table 2 lists the full set of parameters for an allophone. As shown, this includes the parameters for four formants. Note that no bandwidth parameters are included for the fourth formant f4. The bandwidth of the fourth formant is treated as a constant value as it varies little compared with the bandwidth of the other three formants.

10 TABLE 1

DATA FOR ONE ALLOPHONE FORMANT (x)	
Parameter	Description
bbx	frequency at back boundary of allophone
v1x	frequency at time t1
t1x	time of measurement v1
v2x	frequency at time t2
t2x	time of measurement t2
v3x	frequency at time t3
t3x	time of measurement v3
v4x	frequency at time t4
t4x	time of measurement v4
fbx	frequency at forward boundary of allophone
b3x	bandwidth 30 milliseconds after back boundary
b5x	bandwidth 50 percent of the way through the duration of the allophone
b7x	bandwidth 70 percent of the way through the duration of the allophone

30 TABLE 2

DATA FOR ONE ALLOPHONE - FOUR FORMANTS	
FORMANT	Parameters
1	bb1, v11, t11, v21, t21, v31, t31, v41, t41, fb1, b31, b51, b71
2	bb2, v12, t12, v22, t22, v32, t32, v42, t42, fb2, b32, b52, b72
3	bb3, v13, t13, v23, t23, v33, t33, v43, t43, fb3, b33, b53, b73
4	bb4, v14, t14, v24, t24, v34, t34, v44, t44, fb4

45 Data Compression

Using Vector Quantization

To store the parameters listed in Table 2 for a single allophone requires 38 bytes: 8 bytes for the eight forward and back boundary values, 16 bytes for the sixteen intermediate frequency values, 8 bytes for the sixteen intermediate time values (4 bits each), and 6 bytes for the three sets of bandwidth values. Table 3 shows how each measurement value is scaled so as to enable this efficient representation of the data for one allophone. Using more standard, less efficient, representations of the formants would require forty-one or more bytes of data for each allophone.

TABLE 3

FORMANT DATA SCALING		
Parameter(s)	# Bits Used*	Scaling
ALLOPHONE DATA TABLES:		
bb1, fb1	8	value/4
bb2, fb2	8	(value-500)/8
bb3, fb3	8	value/16
bb4, fb4	8	value/16
b3	6	value/8
b5	5	value/12
b7	5	value/12
FX1	10	code book 1 index value
FX2	9	code book 2 index value
FX3	7	code book 3 index value
FX4	6	code book 4 index value
CODE BOOK VALUES:		
v11 thru v41	8	value/4
v12 thru v42	8	(value-500)/8
v13 thru v43	8	value/16
v14 thru v44	8	value/16
t11 thru t44	4	percentage of duration of measured allophone, divided by 2

* number of bits used for each parameter

Note that the amount of data storage needed to store the formant parameters for 72,128 vowel allophones, at 38 bytes per allophone, is 2,740,864 bytes.

35

Formant Code Books

40

One embodiment of the present invention reduces the amount of data storage needed in two ways: (1) by using vector quantization to more efficiently encode the "intermediate" portions of the formants (i.e., v1 through v4 and t1 through t4), and (2) denoting "duplicate" allophones with virtually identical formant parameter sets. This section describes the vector quantization used in the preferred embodiment.

45

Figure 5 depicts a data structure herein called the code book 90 for one formant. Since each allophone is modelled as having four formants, the TTS system uses four code books 90a - 90d, as will be discussed in more detail below.

For the purposes of this example, assume that the code book 90 in Figure 5 has 1000 rows of data. Each entry or row 92 contains the intermediate data values for one allophone formant: v1 through v4 and t1 through t4, as defined in Table 1.

50

Using the code book 90, the data 94 representing one allophone formant is now reduced to forward and back boundary values bb and fb, three bandwidth values b3, b5 and b7, and a pointer 96 to one entry (i.e., row) in the code book. Thus the amount of data storage required to store one allophone formant is now five bytes: one for the pointer 96, two for the boundary values and two for the bandwidth values. For the fourth formant, the amount of storage required is three bytes because no bandwidth data is stored. Without the code book 90, the amount of storage required was ten bytes per formant, and eight for the fourth formant.

55

Thus, if the code book 90 is considered to be a "fixed cost", the amount of storage for each allophone formant is reduced by half through the use of the code book. To show that this is a valid measurement of data compression consider the following. If code books are not used, the amount of data storage required to store the intermediate frequency and time values for 72,128 allophones is 24 bytes per allophone, or a total of 1,731,072 bytes. Four code books, with an average of 1000 entries each, occupy 24,000 bytes. Storing 72,128 allophones, using four one-byte

code book pointers per allophone, requires 288,512 bytes to store the pointers, plus 24,000 bytes for the code books, for a total 312,512 bytes - as compared to 1,731,072 bytes without compression. This represents a compression ratio of about 5.5:1.

5 The next issue is deciding which data values to store in the code book 90 for each formant. In other words, we must choose the 1000 items 92 in the code book 90 wisely so that there will be an appropriate entry for every allophone in the English language.

10 Referring to Figure 6, the four code books 90a - 90d for the four formants f_1 - f_4 are generated as follows. First, the speech of a single, selected person is recorded 100 while speaking each and every vowel allophone in the English (or another selected) language. Next, the recorded speech is digitized and processed to produce a spectrogram for each vowel allophone. Then, a trained technician selects representative formant frequency values from the formant trajectories of each vowel allophone. The result of this process is formant frequency and time data 104 for each of four formants for each of the vowel allophones in the English language. Of course, the process being described here can be performed with data from just a subset of the vowel allophones.

15 It is noted that the TTS system 50 can be made to mimic any selected person, selected dialect, or even a selected cartoon character, simply by recording a person with the desired speech characteristics and then processing the resulting data.

20 There is a well-known technique, called vector quantization, for "mapping" a sequence of continuous or discrete vectors into a smaller representative set of vectors. For a description of how vector quantization works, see Robert M. Gray, "Vector Quantization", IEEE ASSP Magazine, pp. 4-29, April 1984.

25 Suffice it to say that given a set of 288,512 (i.e., $4 * 72,128$) vectors (box 104 in Figure 6) of the form:

$$(v1,t1) (v2,t2) (v3,t3) (v4,t4)$$

25 vector quantization can be used to generate the set of X vectors which produce the minimum "distortion". Given any value of X, such as 4000, the vector quantization process 106 will find the "best" set of vectors. This best set of vectors is called a "code book", because it allows each vector in the original set of vectors 104 to be represented by an "encoded" value - i.e., a pointer to the most similar vector in the code book.

30 Generally, the best set of vectors is one which minimizes a defined value, called the distortion. In the preferred embodiment, the vector quantizer 106 implements a "minimax" method which selects a specified number of code book vectors from the set of all vowel allophone vectors such that the maximum weighted distance from the vectors in the set of vowel allophone vectors to the nearest code book vectors is minimized. The weighted distance between two vectors is computed as the area between the corresponding formant trajectories multiplied by $1/F$, where F is the average of the forward and backward boundary values for the two trajectories. The distance is weighted by $1/F$ to give 35 greater importance to lower frequencies, because lower frequencies are more important than higher ones in human perception of speech. It has been discovered that the minimax method results in higher quality speech than does an alternative method that minimizes the average of the distances from the vowel allophone vectors to their nearest code book vectors. See Eric Dorsey and Jared Bernstein, "Inter-Speaker Comparison of LPC Acoustic Space Using a Minimax Distortion Measure," Proc. IEEE Int'l Conf. Acoustics, Speech and Signal Processing (1981) for a discussion of 40 minimax distortion vector quantization as applied to LPC encoded speech.

45 The vector quantization is performed once on the entire set of vowel allophone vectors representing data for all four formants to generate four formant code books 90a - 90d with a total specified size, such as 4000 rows, for the four code books. In other words, to form code book 90a, the selected vectors that represent formant f_1 are stored in that code book. Similarly, selected vectors for formants f_2 , f_3 and f_4 are stored in code books 90b, 90c and 90d, respectively. The sum $n_1 + n_2 + n_3 + n_4$, where n_x is the number of vectors in the code book for formant f_x , is equal to the total code book size specified when the vector quantization process is performed.

50 In the preferred embodiment, the number of items in each of the code books 90a - 90d is different because the different formants have differing amounts of variability. In general, $n_1 > n_2 > n_3 > n_4$, because use of the $1/F$ weighting factor gives lesser importance to differences between vectors representing higher formants with the result that fewer 55 vectors are selected for the higher formants. This is desirable because each higher formant is less critical to perceived vowel quality than the lower formants. In one version of the preferred embodiment the following values were used: $n_1 = 741$, $n_2 = 451$, $n_3 = 127$ and $n_4 = 81$. However, these values change when the allophone data is changed (e.g., when new allophone data is added). In the preferred embodiment $n_1 + n_2 + n_3 + n_4$ is set to a fixed size, such as 1400 or 4000 (depending on the number of vectors being quantized), and the quantizer sets the individual sizes to minimize the overall weighted distortion.

Once all of the code books have been generated, vector quantization is no longer used. Thus the completed TTS system need not incorporate a vector quantization capability. In the completed TTS system, each allophone is "encoded" using the four formant code books 90a - 90d with the parameters shown in Table 4.

TABLE 4

PARAMETERS FOR ONE ALLOPHONE	
Parameter(s)	Description
FX1 - FX4	indices to entries in formant code books 1, 2, 3 and 4
bb1 - bb4	frequency at back boundary of allophone for formants 1 - 4
fb1 - fb4	frequency at forward boundary of allophone for formants 1 - 4
b31 - b33	bandwidth 30 milliseconds after back boundary for formants 1 - 3
b51 - b53	bandwidth 50 percent of the way through the duration of the allophone, for formants 1 - 3
b71 - b73	bandwidth 70 percent of the way through the duration of the allophone, for formants 1 - 3
LLRRRx	index into LLRR Context Table
LLRRd	index into LLRR Allophone Data Table for corresponding vowel phoneme

It should be noted that in the preferred embodiment, the formant data in the code books 90a - 90d is derived from the speech of a single person, though the data for any particular vowel allophone may represent the most representative of several enunciations of the vowel allophone. This is different from most TTS synthesis systems and methods in which the formant and bandwidth data stored to represent phonemes is data which represents the "average" speech of a number of different persons. The inventors have found that the averaging of speech data from a number of persons tends to average out the tonal qualities which are associated with natural speech, and thus results in artificial sounding synthetic speech.

Generating Vowel Allophones

When converting text to speech vowel phonemes are converted into vowel allophones using the process shown in Figures 7 through 10. It is to be noted that the process of converting vowel phonemes is performed between boxes 30 and 40 in the flow diagram of Figure 1. Thus, at the beginning of this process, the phonemes preceding and following the vowel phoneme to be converted (the currently "selected" vowel phoneme) are known.

For the purposes of this discussion, it should be understood that the term "vowel allophone" refers to the particular pronunciation of a vowel phoneme as determined by its neighboring phonemes. As explained below, there is conceptually a distinct allophone for every PVP context of the vowel phoneme V. However, some allophones are perceptually indistinguishable from others. For this reason, some vowel allophones are labelled "duplicate" allophones. To save on memory storage, the formant data representing such duplicate allophones is not repeated.

Many vowels are diphthongs, gliding speech sounds that start with the acoustic characteristics of one vowel and move toward having those of another. The second part of a diphthong is called an "offglide". There are just a few, common offglides, so vowels fall into a few groups that have a common offglide, and therefore a common effect on a following phoneme. This has enabled the inventors to group preceding and following vowels into a few categories and to simplify the system to store and process 1156 (i.e., 34 x 34) CVC (i.e., consonant-vowel-consonant) contexts plus several CVV (i.e., consonant-vowel-vowel), VVC (i.e., vowel-vowel-consonant) and VVV (vowel-vowel-vowel) contexts for each vowel phoneme instead of all 3136 (i.e., 56 x 56) PVP (phoneme-vowel-phoneme) contexts for each vowel.

Referring to Figure 7A, the first step of the vowel phoneme conversion process is to determine the context of the vowel phoneme. The identity of the most appropriate vowel allophone to be used is initially determined by the identity of the phonemes preceding and following selected vowel phoneme.

Figure 7A shows a context index calculator 110. The input data to the context index calculator 110 are the phonemes P1 and P2 preceding and following the selected vowel phoneme V. Initially we will assume that the neighboring phonemes are consonant phonemes. Of course, sometimes one of both of the neighboring phonemes are vowels, but we will deal with those cases separately.

The Phoneme Index Table 112 converts any phoneme into an index value between 0 and 33, i.e., one of 34 distinct values. In the preferred embodiment, there are 33 distinct consonant phonemes plus one for silence. Thus Phoneme

Index Table 112 generates a unique value for each consonant phoneme, including the silence phoneme.

The Phoneme Index Table 112 is used to generate two index values I1 and I2, corresponding to the identities of the two neighboring phonemes P1 and P2, respectively. The context index calculator 110 then generates a CVC index value:

5

$$\text{CVC Index} = I2 + 34*I1$$

which uniquely identifies the "context" of a vowel phoneme - i.e., the preceding and following consonant phonemes.

10 In most cases, the CVC Index value can be used to correctly identify the vowel allophone associated with the vowel V.

When one of the neighboring phonemes is a vowel, the inventors have found that, for the purposes of selecting the most appropriate allophone, the following substitution process can be used.

TABLE 5

ALLOPHONE SUBSTITUTION TABLE FOR C-V1-V2 and V1-V2-C CONTEXTS	
V1	REPLACE OUTER VOWEL WITH CONSONANT INDEX FOR:
/ej/, /ij/, /ai/, or /ɔi/	/j/
/ou/, /juw/, /uw/, /ɔ/, or /au/	/w/
/ɔ/, /ir/, /er/, /ur/, /ɔr/, or /ar/	/r/
/ə/, /a/, /ʌ/, /æ/, /ɛ/, /ɪ/, /ʊ/, or /ʊ/	/?/

25 The PUP context is relabelled C-V1-V2, or V1-V2-C, as appropriate. To synthesize the inner vowel (V1 in the first case, V2 in the second), use the substitution values shown in Table 5 (in which phonemes are denoted using standard IPA symbols) so that a consonant is substituted for the outer vowel. Then the CVC index is computed, as explained above.

20 To implement the vowel substitutions shown in Table 5, the Phoneme Index Table 112 includes entries for the 23 vowel phonemes. The entries in the Phoneme Index Table 112 for vowel phonemes are set equal to the values for the substitute consonant phonemes specified in Table 5. Thus, the context of any and all vowel phonemes is computed simply by looking up the index values for the neighboring phonemes (regardless of whether they are consonants or vowels) and then using the CVC index formula shown above.

30 It is to be noted that the "substitution" represented in Table 5 is used solely for the purpose of generating a CVC index value to represent the context of the selected vowel phoneme V. The original "outer vowel" is used when synthesizing the outer vowel.

35 Thus, at this point, whether the neighboring phonemes are consonants or vowels, we have a CVC index value representing the context of a selected vowel phoneme V.

40 Referring to Figure 7B, the formant parameters for a selected vowel phoneme V are generated as follows. There are 23 vowel phoneme-to-allophone decoders 120, one for each of the 23 vowel phonemes. As will be described in more detail, each vowel phoneme-to-allophone decoder 120 stores encoded data representing all of the vowel allophones for the corresponding vowel phoneme.

45 Whenever a vowel phoneme is encountered in the string of phonemes that is being synthesized, the data for the corresponding allophone is generated as follows. First, the CVC index for the context of the vowel phoneme is calculated, as described above with reference to Figure 7A. Then, the CVC index is sent by a software multiplexer 122 to the allophone decoder 120 for the corresponding vowel phoneme V.

The selected allophone decoder 120 outputs four code book index values FX1 - FX4, as well as a set of formant data values FD which will be described below. The allophone decoder 120 is shown in more detail in Figure 7C. The code books 90a - 90d output formant data FDC representing the central portions of the four speech formants for the selected vowel allophone.

50 The combined outputs FD and FDC are sent to a parameter stream generator 124, which outputs new formant values to the formant synthesizer 62 (shown in Figure 2) once every 10 milliseconds for the duration of the allophone, thereby synthesizing the selected allophone. More generally, the parameter stream generator 124 continuously outputs formant data every 10 milliseconds to the formant synthesizer, with the formant data representing the stream of phonemes and/or allophones that are selected by earlier portions of the TTS conversion process.

55 Figure 7C shows one vowel phoneme-to-allophone decoder 120. As explained above, there are 23 such decoders, one for each of the 23 vowel phonemes in the preferred embodiment. Thus the data stored in the decoder 120 represents the allophones for one selected vowel phoneme.

The data representing all of the allophones associated with one vowel phoneme V is stored in a table called the

Allophone Data Table 130.

Referring to Figure 8, each Allophone Data Table 130 contains separate records or entries 132 for each of a number of unique vowel allophones. Each record 132 in the Allophone Data Table 130 contains the set of data listed in Table 3, as described above. In particular, the record 132 for any one allophone contains four code book indices FX1 - FX4, representing the center portions of the four formants f1 - f4 for the allophone, four values bb1 - bb4 representing the back boundary values of the four formants, four values fb1 - fb4 representing the forward boundary values of the four formants, nine bandwidth values b31 - b73 representing the bandwidths of the three lower formants f1 - f3 (as shown in Figure 3), and a value called LLRR which will be described below.

The data values in the record 132 are scaled using the scaling and compression factors listed in Table 3. As a result, each record 132 occupies 19 bytes in the preferred embodiment.

The Allophone Data Table 130 has two portions: one portion 134 for allophones identified by the PVP context (i.e., the CVC index value) of the Vowel V, and a smaller portion 136 for the allophones identified by the expanded context LCVC or CVCR of the vowel V as will be explained in more detail below. The smaller portion 136, called the Extended Allophone Data Table, contains up to 16 records, each having the same formant as the records in the rest of the table 130.

While there are 1156 possible CVC contexts for each vowel phoneme V, the inventors have further reduced memory requirements by selecting a number of "distinct allophones" which sound sufficiently distinct to require storage. The number of distinct allophones represented in the preferred embodiment is around 10,000 (less than half the number of CVC contexts), with the exact number depending on the methodology used to select them. Thus many vowel allophones are perceptually similar and can be considered to be "duplicate" allophones. It is noted that the selection of distinct allophones is inherently subjective, since it based on judgments by human technicians.

Storing formant data for 26,588 allophones would require 505,172 bytes of storage (excluding the storage required for the code books 90a - 90d). On the other hand, storing formant data for only the 10,000 or so distinct allophones requires about 190,000 bytes of storage - which is a significant savings of memory storage for low cost TTS systems.

As a result, only the distinct vowel allophones for a selected phoneme V are stored in each Allophone Data Table 130.

Referring to Figure 7C, the purpose of the Allophone Context Table 140, Duplicate Context Table 144, and LLRR Table 148 is to enable the use of a compact Allophone Data Table 130 which stores data only for distinct allophones. These additional tables 140, 144 and 148 are used to convert the initial CVC index value into a pointer to the appropriate record in the Allophone Data Table 130.

Figure 9 shows an Allophone Context Table 140, for one phoneme V. The purpose of the Allophone Context Table 140 is to convert a CVC index value (calculated by the indexing mechanism shown in Figure 7A) into a Context Index CI.

Each of the 23 Allophone Context Tables 140 contains a single Mask Bit, Mask(i), for each of the 1156 CVC contexts for a vowel phoneme V. Distinct vowel allophones are denoted with a Mask Bit 142 equal to 1, and "duplicate" vowel allophones which are perceptually similar to one of the other vowel allophones are denoted with a Mask Bit of 0. Nonexistent allophones (i.e., CVC contexts not used in the English language) are also denoted with a Mask Bit equal to 0.

To find the CI index value for any particular vowel allophone, the Mask value Mask(CVC Index) is inspected. If the Mask Bit value is equal to 1, the value of CI is computed as the sum of all the Mask Bits for CVC Index values less than or equal to the selected CVC Index value:

$$CI(N) = \sum_{i=0}^N \text{Mask}(i)$$

where N is equal to the CVC Index value that is being converted into a CI value.

The number of unique vowel allophones for the selected vowel phoneme is CIMAX(V), which is also equal to CI for the largest CVC index with a nonzero Mask Bit. CIMAX(V) is furthermore equal to the number of records 132 in the main portion 134 of the Allophone Data Table 130. Referring to Figure 8, the number of entries 132 in the Allophone Data Table 130 is CIMAX(V) + 16, for reasons which will be explained below.

If the selected Mask Bit 142 equals 0, the selected allophone is a "duplicate", and a substitute CVC index value is obtained from the Duplicate Context Table 144. The substitute CVC index value is guaranteed to have a Mask Bit equal to 1, and is used to compute a new CI index value as described above.

More particularly, to find the CI value for a particular "duplicate" allophone, the synthesizer looks through the records 146 of the Duplicate Context Table 144 for the CVC index value of the duplicate allophone. When the CVC index value is found, the new CVC value in the same record replaces the original CVC index value, and the CI computation process is restarted.

As shown in Figure 9, the Duplicate Context Table 144 comprises a list of "old" or original CVC Index Values and

corresponding "new CVC" values, with two bytes being used to represent each CVC value. In other words, the Table 144 comprises a set of four byte records 146, each of which contains a pair of corresponding CVC Index and "new CVC" values. The only "old" CVC Index values included in the Duplicate Context Table 144 are those for existent allophones which have a Mask Bit value of 0 in the Allophone Context Table 140. Thus the Duplicate Context Table 144 will typically contain many fewer records 146 than there are Mask Bits 142 with values of zero. In the preferred embodiment, the number of entries in the Duplicate Context Table 144 varies from 24 to 111, depending on the vowel phoneme V.

Should the selected CVC value not be found in the Duplicate Context Table, this would mean that a previously unknown allophone context has been encountered. In this case, the TTS synthesizer synthesizes the allophone using a standard "default" context for all allophones. In an alternate embodiment, such allophones could be synthesized using the "synthesis by rule" methodology previously used in Speech Plus Prose 2000 product (described above with reference to Figure 1).

In another embodiment of the invention, the Duplicate Context Table 144 stores the CI value for each duplicate allophone. Since the CI value occupies the same amount of storage space as a replacement CVC value, the alternate embodiment avoids the computation of CI values for those allophones which are "duplicate" allophones.

In yet another alternate embodiment of the invention, the Allophone Context Table 140 (for one vowel V) comprises a table of two byte index values CI, with one CI value for each of the 1156 possible CVC index values. By eliminating the Mask Table 150, the alternate embodiment occupies about 2000 bytes of extra storage per vowel phoneme V, but reduces the computation time for calculating CI.

Referring to Figure 7C, we now have a CI index value which points to one record in the Allophone Data Table 130. As mentioned above, the data in each record 132 of the Allophone Data Table 130 includes an entry called LLRR. LLRR actually has two components: LLRRx (the low-order four bits) and LLRRd (the high-order four bits).

LCVC and CVCR Contexts

In a relatively small number of cases, the selection of the proper vowel allophone depends not just on the immediately neighboring phonemes, but also on the phoneme just to the left or to the right of these neighboring phonemes. The "expanded" context of selected vowel phoneme can be labelled:

LCVC or CVCR.

Thus there are multiple allophones for a small number of CVC contexts. The inventors have found that, for any one CVC context, there is at most one LCVC or CVCR context which has a distinct enunciation of the vowel allophone V. As a result, a relatively small LLRR Context Table 148 and a similarly small Extended Allophone Data Table 136 can be used to represent and store the formant data for these allophones.

The LLRRx value in each Allophone Data Table record denotes whether there is more than one allophone for the selected CVC context, and thus whether the "expanded" LCVC or CVCR context of the allophone must be considered. If LLRRx is equal to zero, the allophone data specified by the previously calculated value of CI is used. If LLRRx is not equal to zero, then an additional computation is needed.

Referring to Figure 10, there is an LLRR Context Table 148 for each vowel phoneme V. The Table 148 contains fifteen entries or records, each of which identifies an "extended" context. More particularly, the Table 148 can denote up to fifteen Left or Right Phonemes which identify an extended LCVC or CVCR context.

Each LLRR Context Table record has two values: LRI and CC. The value of LLRRx determines which entry in the Table 148 is to be used. Note that there is no entry for LLRRx = 0 because a value of zero indicates that the expanded context need not be considered.

CC denotes a phoneme value, and LRI is a "left or right" indicator. When LRI is equal to 0, the phoneme to the left of the CVC context is compared with the phoneme denoted by CC; when LRI is equal to 1, the phoneme to the right of the CVC context is compared with the CC phoneme. Only if the selected left or right phoneme matches the CC phoneme is a "new LLRR CI value" calculated.

If the selected left or right phoneme does not match the CC phoneme, then the data pointed to by CI is the data used to generate the allophone. If there is a match, however, the LLRRd value acts as a pointer to a record in the extended portion 136 of the Allophone Data Table 130 shown in Figure 8. In effect, the CI value is replaced with a value of

$$\text{CIMAX}(V) + \text{LLRRd}$$

where CIMAX(V) is the number of records in the main portion 134 of the Allophone Data Table 130.

While there are only sixteen possible values of LLRRd in the preferred embodiment, in alternate embodiments a full byte could be used to represent LLRRd, allowing for a much larger number of extended context allophones. Note that there is not a one to one correspondence between the entries in the LLRR Table 148 and the Extended Allophone

Data Table 136. In fact, there can be several Extended Allophone Data Table entries for a single LLRR Table entry because one LLRR Table entry can define the context of several allophones.

Allophone Synthesis Method

5 Referring once again to Figure 7C, the process for synthesizing a particular vowel phoneme V is as follows. First a CVC index value is computed by the context index calculator 110. Then, using the allophone decoder 120 for the selected vowel phoneme V, a CI index value is computed using the Allophone Context Table 140 and Duplicate Context Table 144. The CI index value points to a record in the Allophone Data Table 130, which contains formant data for the 10 allophone. However, if the LLRR value in the selected Allophone Data record has a value of LLRR_x ≠ 0, and the expanded context LCVC or CVCR matches the specified value in the LLRR Table 148, a new CI value replaces the old one and a new record of data in the Allophone Data Table 130 is used.

15 The data record 132 of the Allophone Data Table 130 pointed to by CI includes four pointers FX1 - FX4 to records in the four formant code books 90a - 90d. The data record 132 also includes back boundary and forward boundary values for the four formants, and a sequence of three bandwidth values for each of the first three formants. The formant parameters representing the four formant frequency trajectories for the vowel allophone include the data values from the four selected code book records as well as the data values in the selected Allophone Data Table record.

20 These formant parameters are then processed by a parameter stream generator 124. This generator 124 interpolates between the selected formant values to compute dynamically changing formant values at 10 millisecond intervals from the start of the vowel to its end. For each formant, quadratic smoothing is used from the back boundary at the start of the vowel to the first "target" value retrieved from the code book. Linear smoothing is performed between the four target values retrieved from the code book, and also between the fourth code book value and the forward boundary value at the end of the vowel.

25 Most contexts require smoothing of the formants backward into the preceding consonant in order to assure a continuous formant track. To do this, interpolation is done from the vowel's back boundary value to a formant value in the preceding consonant. Consonants for which this is not done are those where a discontinuity is desired in formants f2, f3 and f4, namely the nasal consonants (m, n and ng) and stop consonants (p, t, k, b, d, g).

30 For each formant, the bandwidth is linearly smoothed from the last bandwidth value of the preceding phoneme to the 30 ms bandwidth value b3x, then to the midpoint bandwidth value b5x, then to the 75% value b7x, and then to the boundary of the next phoneme.

Alternate Embodiments

35 While the present invention has been described with reference to a few specific embodiments, the description is illustrative of the invention and is not to be construed as limiting the invention.

In particular, it is noted that the data compression methods used in the preferred embodiment are dictated by the need to store all the vowel allophone data in a space of 256k bytes or less. If the storage space limits are relaxed, because of relaxed cost criteria or reduced memory costs, a number of simplifications of the data structures well known to those skilled in the art could be employed.

40 For instance, as noted above, the allophone context table 140 and duplicate context table 144 could be combined and simplified at a cost of around 45k bytes. At a cost of approximately 256k, formant data can be stored for every CVC context, thereby eliminating the need for the Allophone Context Table 140 and Duplicate Context Table 144 altogether.

45 In other alternate embodiments, bandwidth values could be stored in code books much as the formant values are stored in the preferred embodiment. Similarly, code books could be used to store formant parameter vectors that include the backward and forward formant boundary values (instead of the above described code books, which store vectors that include only the intermediate formant parameters). These alternate embodiments would increase the amount of data compression obtained from the use of code books, but would degrade the quality of the synthesized allophones.

50 It is also noted that each TTS system can store allophone data representative of the pronunciation of a selected individual, a selected dialect, a selected cartoon character, or a language other than English. The only difference between these embodiments of the present invention's vowel allophone production system is the allophone data stored in the system. In still other embodiments in which there is more memory available for allophone storage, multiple sets of allophone data could be stored so that a single TTS system could generate synthetic speech which mimics several different persons or dialects.

55 Finally, it is noted that in an alternate embodiment of the present invention vowel allophones could be stored using speech parameters that are based on a different representation of human speech than the formant parameters described above. It is well known to those skilled in the art that there are several alternate methods of representing synthetic speech using speech parameters other than formant parameters. The most widely used of these other methods is

known as LPC (linear predictive coding) encoded speech.

Referring to Figure 11, in an alternate embodiment of the invention each distinct vowel allophone is represented by a set of stored LPC encoded data. Note that Figure 11 is the same as Figure 7C, except for the data and code book tables. The LPC data for each vowel allophone is a set of parameters which can be considered to be a vector. Synthetic speech is generated from LPC parameters by processing the LPC parameters with a digital signal processor (i.e., a digital filter network). While the digital signal processors used with LPC parameters are different than the digital signal processors used with formant parameters, both types of digital signal processors are well known in the prior art and can be considered to be analogous for the purposes of the present invention.

Since the LPC parameters for each vowel allophone is a vector, the amount of storage required to represent these vectors can be greatly reduced using the vector quantization scheme described above. In particular, the intermediate portions of the LPC vectors for all the vowel allophones can be processed by a minimax distortion vector quantization process, as described above, to produce the best set of N vectors (e.g., 4000 LPC vectors) for representing the intermediate portions of the LPC vectors. The resulting N vectors would be stored in a single parameter code book 152.

The LPC Allophone Data Table 150 will store forward and back LPC boundary values, bandwidth values, LLRR, and a single index into the parameter code book 152.

The methodology for selecting vowel allophones and retrieving the data representing a selected vowel allophone is unchanged from the preferred embodiment, except that now there is only one code book entry that is retrieved (instead of four). The parameters selected from the Allophone Data Table 150 and the parameter code book 152 are sent to the parameter stream generator 124 for inclusion in the stream of data sent to the synthesizer's digital signal processor.

In yet other embodiments of the present invention, other methods of representing vowel allophones with speech parameters can be used. Several such alternate methods are known to the prior art, and new parameter representations of speech may be developed in the future.

In all such alternate embodiments, the primary differences from the preferred embodiment would be in the vowel allophone data stored, and in the apparatus used to convert the vowel allophone data into synthetic speech. The number of code books used to compress the vowel allophone parameters will vary depending on the nature of parameter representation being used. Nevertheless, the system architecture shown in Figure 11 can be applied to all of these embodiments because the basic methodology for selecting vowel allophones and retrieving the data representing a selected vowel allophone is unchanged.

Claims

1. A text-to-speech synthesis system, comprising:

35 text conversion means (20, 22, 24) for converting a specified text string into a corresponding string of consonant and vowel phonemes (25), each said phoneme being selected from a predefined set of phonemes including a multiplicity of consonant phonemes and a multiplicity of vowel phonemes; parameter generating means (40) for generating speech parameters corresponding to said string of phonemes (25); and speech synthesizing means (42) for generating a speech waveform corresponding to the speech parameters generated by said parameter generating means;

45 characterised by:

vowel allophone storage means (90, 130) storing a multiplicity of predefined vowel allophones, each vowel allophone being represented by a set of speech parameters; said vowel allophones including allophones for a multiplicity of vowel phonemes;

50 vowel phoneme to allophone conversion means (120), coupled to said text conversion means (20, 22, 24) and said vowel allophone storage means, for computing a phoneme context value for each of at least a subset of said vowel phonemes in said string of phonemes (25), said phoneme context value comprising a function of the phonemes in said string of phonemes (25) which precede and follow said vowel phoneme, and for then assigning to said vowel phoneme a selected one of said predefined vowel allophones corresponding to said computed phoneme context value;

55 said parameter generating means (40) including means for generating speech parameters for said assigned vowel allophones.

2. The text-to-speech conversion system of claim 1, further characterised by:

context table means (140) for assigning one of said vowel allophones to every vowel phoneme context LVR, where V represents any vowel phoneme selected from said multiplicity of vowel phonemes, L represents any consonant phoneme immediately preceding said vowel phoneme V selected from said predefined set of phonemes, and R represents any consonant phoneme immediately following said vowel phoneme V selected from said predefined set of phonemes; said context table means (140) including a distinct entry for every phoneme context LVR denoting which of said vowel allophones is assigned to each said phoneme context LVR; and

said vowel phoneme to allophone conversion means (120) including allophone selection means coupled to said context table means (140) for selecting one of said multiplicity of vowel allophones for each of at least a subset of said vowel phonemes in said string of phonemes (25), said allophone selection means including context indexing means 110 for determining the phonemes in said string which immediately precede and follow said vowel phoneme in said string of phonemes, and table lookup means for assigning to said vowel phoneme the vowel allophone denoted in said context table means (140) for said vowel phoneme in the context of said preceding and following phonemes.

15 3. The text-to-speech conversion system of claim 1 or claim 2, further characterised by:

said vowel allophone storage means (90, 130) including:
 speech storage means for storing the speech parameters for each said vowel allophone; said speech storage means including code book means (90) for storing a multiplicity of sets of speech parameters; and
 20 allophone means (130) for denoting, for each said vowel allophone, one of said multiplicity of sets of speech parameters in said code book means (90).

25 4. The text-to-speech conversion system of claim 2, further characterised by:

said context indexing (110) means including vowel substitution means (112) for use when a vowel phoneme V₁ in said string of phonemes (25) is immediately preceded or followed by a vowel phoneme, said vowel substitution means (112) including means for selecting an entry in said context table means (140) to use for assigning one of said vowel allophones to said vowel phoneme V₁.

30 5. The text-to-speech conversion system of claim 2, further characterised by:

said context indexing means (110) including vowel substitution means (112) for use when a vowel phoneme V₁ in said string of phonemes (25) occurs in a phoneme context CV₁V₂ or V₂V₁C, where C is a consonant phoneme and V₁ is a vowel phoneme neighboring said vowel phoneme V₁, said vowel substitution means (112) including means for selecting one of said phoneme contexts LVR which is phonetically equivalent to said phoneme context CV₁V₂ or V₂V₁C; said table lookup means including means for assigning to said vowel phoneme V₁ the vowel allophone denoted in said context table means (140) for said phonetically equivalent phoneme context LVR.

35 6. The text-to-speech conversion system of any one of claims 1 to 5, characterised in that:
 said speech parameters are formant parameters.

40 7. The text-to-speech conversion system of claim 6, characterised in that:

the number of sets of formant parameters stored in said code book means (90) is much less than the number of vowel allophones stored by said vowel allophone storage means (90, 130); the sets of formant parameters stored in said code book means (90) being selected from sets of formant parameters representing substantially all of said vowel allophones using a minimax distortion vector quantization process.

45 8. The text-to-speech conversion system of any one of claims 1 to 6, further characterised by:

each vowel allophone in said vowel allophone storage means including a set of back and forward boundary parameters representative of speech formants at the boundaries of the allophone, and a set of intermediate parameters representative of speech formants between the back and forward boundaries of the allophone.

50 9. The text-to-speech conversion system of claim 8, further characterised by:

each said set of intermediate parameters in said code book means (90) representing the intermediate trajectory of one formant for a vowel allophone; said allophone storage means (90, 130) including means for denoting at least three of said sets of intermediate formant parameters; whereby said vowel allophones comprise the formant parameters for at least three formants.

55 10. The text-to-speech conversion system of any one of claims 1 to 9, further characterised by:

said vowel allophone storage means (90, 130) including means for storing vowel allophones as pronounced by a selected individual so that said text-to-speech conversion system produces synthetic speech which mimics said selected individual speaking.

5 11. The text-to-speech conversion system of any one of claims 1 to 9, further characterised by:
 said vowel allophone storage means (90, 130) including means for storing vowel allophones as pronounced by an individual speaking a selected dialect so that said text-to-speech conversion system produces synthetic speech which mimics said selected dialect.

10 12. The text-to-speech conversion system of any one of claims 1 to 9, further characterised by:
 said vowel allophone storage means including means for storing vowel allophones as pronounced by a specified cartoon character so that said text-to-speech conversion system produces synthetic speech which mimics said selected cartoon character.

15 13. The text-to-speech conversion system of any one of claims 1 to 9, further characterised by:
 said vowel allophone storage means (90, 130) including means for storing vowel allophones as pronounced by a plurality of selected individuals so that said text-to-speech conversion system produces synthetic speech which mimics a plurality of selected individuals.

20 14. A method of converting text strings into synthetic speech, the steps comprising:
 defining a set of phonemes, including a multiplicity of consonant phonemes and a multiplicity of vowel phonemes;
 converting a specified text string into a corresponding string of phonemes (25), said string of phonemes including consonant and vowel phonemes, each said phoneme being selected from said defined set of phonemes; and
 converting said string of phonemes (25) into speech parameters and then generating an audio waveform corresponding to said speech parameters;

25 characterised by:
 storing a multiplicity of predefined vowel allophones, each vowel allophone being represented by a set of speech parameters;
 for each of at least a subset of said vowel phonemes in said string of phonemes (25), computing a phoneme context value for said vowel phoneme as a function of the phonemes in said string of phonemes which precede and follow said vowel phoneme, and then assigning to said vowel phoneme a selected one of said predefined vowel allophones corresponding to said computed phoneme context value; and
 said converting step including converting said assigned vowel allophones into speech parameters which are then used to generate an audio waveform corresponding to said speech parameters.

30 15. The method of claim 14, further characterised by:
 storing a multiplicity of predefined vowel allophones, each vowel allophone being represented by a set of speech parameters; denoting in a data structure an assigned one of said vowel allophones for every phoneme context LVR, where V represents any vowel phoneme selected from at least a subset of said multiplicity of vowel phonemes, L represents any consonant phoneme immediately preceding said vowel phoneme V selected from said predefined set of phonemes, and R represents any consonant phoneme immediately following said vowel phoneme V selected from said predefined set of phonemes; said data structure containing a distinct allophone assignment entry for each said phoneme context LVR; and
 for each vowel phoneme in at least a subset of said vowel phonemes in said string of phonemes (25), determining the phonemes in said string which immediately precede and follow said vowel phoneme in said string of phonemes, and then assigning said vowel phoneme the vowel allophone denoted in said data structure for said vowel phoneme in the context of said preceding and following phonemes.

35 16. The method of claim 14, further characterised by:
 said storing step including providing code book means (90) for storing a multiplicity of sets of speech parameters, and providing allophone means (130) for denoting, for each said vowel allophone, one of said multiplicity of sets of speech parameters in said code book means (90).

17. The method of claim 16, characterised in that:

the number of sets of speech parameters stored in said code book means (90) is much less than said pre-defined multiplicity of vowel allophones; the sets of speech parameters stored in said code book means (90) being selected from sets of speech parameters representing substantially all of said vowel allophones using a minimax distortion vector quantization process.

18. The method of any one of claims 14 to 17, further characterised in that:

said storing step storing vowel allophones as pronounced by a selected individual so that said method produces synthetic speech which mimics said selected individual speaking.

19. The method of any one of claims 14 to 18, characterised in that:

said speech parameters are formant parameters.

20. The method of claim 19, further characterised by:

said storing step including providing code book means (90) for storing a multiplicity of sets of formant parameters, and allophone means (130) for denoting, for each said vowel allophone, one of said multiplicity of sets of formant parameters in said code book means (90).

21. The method of claim 19, characterised in that:

the number of sets of formant parameters stored in said code book means (90) is much less than said pre-defined multiplicity of vowel allophones; the sets of formant parameters stored in said code book means (90) being selected from sets of formant parameters representing substantially all of said vowel allophones using a minimax distortion vector quantization process.

22. The method of claim 18, further characterised by:

said storing step including storing vowel allophones as pronounced by a selected individual so that said method produces synthetic speech which mimics said selected individual speaking.

30 **Patentansprüche**

1. System zur Synthese von gesprochener Sprache aus Text mit den folgenden Elementen:

35 Textumsetzeinrichtungen (20, 22, 24) zum Umsetzen eines bestimmten Textstrings in einen entsprechenden String von Konsonanten- und Vokalphonemen (25), wobei das jeweilige Phonem aus einer vorbestimmten Gruppe von Phonemen ausgewählt wird, die aus vielen Konsonantphonemen und vielen Vokalphonemen besteht;
einer Parametererzeugungseinrichtung (40) zum Erzeugen von Sprechparametern, die dem Phonemstring (25) entsprechen; und
40 einer Sprechsprachsyntheseeinrichtung (42) zum Erzeugen von Sprachwellen, die den von Parametererzeugungseinrichtung erzeugten Sprechparametern entsprechen;

gekennzeichnet durch:

45 Vokalallophonspeichereinrichtungen (90, 130), die eine Vielzahl vorbestimmter Vokalallophone speichern, wobei jedes Vokalallophon durch einen Satz von Sprechparametern repräsentiert wird; wobei die Vokalallophone Allophone für eine Vielzahl von Vokalphonemen enthalten;
eine Vokalphonem-Allophon-Umsetzeinrichtung (120), die mit den Textumsetzeinrichtungen (20, 22, 24) und Vokalallophonspeichereinrichtungen in Verbindung steht, die dazu dient, einen Phonemkontextwert für jedes aus mindestens einer Untergruppe der Vokalphoneme im Phonemstring (25) zu berechnen, wobei der Phonemkontextwert eine Funktion der Phoneme in dem Phonemstring (25) aufweist, die dem Vokalphonem vorangehen und folgen, und dann dazu, dem Vokalphonem ein ausgewähltes der vorbestimmten Vokalallophone zuzuordnen, das dem berechneten Phonemkontextwert entspricht;
50 wobei die Parametererstellungseinrichtung (40) eine Einrichtung zum Erzeugen von Sprechparametern für die zugeteilten Vokalallophone aufweist.

55 2. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 1, weiter gekennzeichnet

durch eine Kontexttabelleneinrichtung (140) die dazu dient, einem jeden Vokalphonemkontext LVR eines der Vokalallophone zuzuordnen, wobei V ein aus der Vielzahl von Vokalphonemen ausgewähltes Vokalphonem repräsentiert, L ein dem Vokalphonem V unmittelbar vorausgehendes Konsonantphonem, das aus der vorbestimmten Gruppe von Phonemen ausgewählt ist, repräsentiert und R ein dem Vokalphonem V unmittelbar folgendes Konsonantphonem, das aus der vorbestimmten Gruppe von Phonemen ausgewählt ist, repräsentiert; wobei die Kontexttabelleneinrichtung (140) für jeden Phonemkontext LVR einen eigenen Eintrag hat, in dem festgelegt wird, welche der Vokalallophone dem jeweiligen Phonemkontext LVR zugeordnet werden; und dadurch, daß die Vokalphonem-Allophon-Umsetzeinrichtung (120) eine mit der Kontexttabelleneinrichtung (140) verbundene Allophonauswahleinrichtung zum Auswählen eines der vielen Vokalallophone für jedes von mindestens einer Untergruppe der Vokalphoneme in dem Phonemstring (25), wobei die Allophonauswahleinrichtung eine Kontextindizierungseinrichtung (110) zum Bestimmen derjenigen Phoneme in dem String aufweist, die dem Vokalphonem in dem Phonemstring unmittelbar vorausgehen bzw. folgen, und eine Tabellensucheinrichtung, die dazu dient, dem Vokalphonem das in der Kontexttabelleneinrichtung (140) aufgeführte Vokalallophon für das Vokalphonem im Kontext des vorhergehenden und nachfolgenden Phonems zuzuordnen.

3. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 1 oder 2, weiter dadurch gekennzeichnet daß

die Vokalallophonspeichereinrichtungen (90, 130) die folgenden Elemente aufweisen:
 eine Sprachspeichereinrichtung zum Speichern der Sprechparameter für das jeweilige Vokalallophon; wobei die Sprachspeichereinrichtung eine Codebucheinrichtung (90) zum Speichern vieler Sprechparametersätze aufweist; und
 eine Allophoneinrichtung (130) zum Festlegen eines der vielen Sprechparametersätze in der Codebucheinrichtung (90) für das jeweilige Vokalallophon.

4. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 2, weiter dadurch gekennzeichnet daß die Kontextindizierungseinrichtung (110) eine Vokalsubstituierungseinrichtung (112) aufweist, die dann eingesetzt wird, wenn einem Vokalphonem V₁ in dem Phonemstring (25) ein Vokalphonem unmittelbar vorangeht bzw. folgt, wobei die Vokalsubstituierungseinrichtung (112) eine Einrichtung zum Auswählen eines Eintrags in der Kontexttabelleneinrichtung (140) aufweist, die dem Vokalphonem V₁ eines der Vokalallophone zuteilt.

5. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 2, weiter dadurch gekennzeichnet daß die Kontextindizierungseinrichtung (110) eine Vokalsubstituierungseinrichtung (112) aufweist, die dann eingesetzt wird, wenn ein Vokalphonem V₁ in dem Phonemstring (25) in einem Phonemkontext CV₁V₂ oder V₂V₁C vorkommt, wobei C ein Konsonantphonem und V₂ ein neben dem Vokalphonem V₁ vorkommendes Vokalphonem ist, wobei die Vokalsubstituierungseinrichtung (112) eine Einrichtung zum Auswählen eines der Phonemkontakte LVR aufweist, der phonetisch mit dem Phonemkontext CV₁V₂ oder V₂V₁C äquivalent ist; wobei die Tabellensucheinrichtung eine Einrichtung, die dazu dient dem Vokalphonem das in der Kontexttabelleneinrichtung (140) aufgeführte Vokalallophon V₁ für den phonetisch äquivalenten Phonmenkontext LVR zuzuordnen, aufweist.

6. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 5, dadurch gekennzeichnet daß die Sprechparameter Formantenparameter sind.

7. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 6, dadurch gekennzeichnet daß die Anzahl in der Codebucheinrichtung (90) gespeicherter Formantenparametersätze weit geringer ist als die Anzahl in den Vokalallophonspeichereinrichtungen (90, 130) gespeicherter Vokalallophone; wobei die Sätze in der Codebucheinrichtung (90) gespeicherter Formantentparameter unter Verwendung eines Minimaxverzerrungsvektorquantisierungsverfahrens aus Sätzen von Formantenparametern ausgewählt werden, die im wesentlichen alle Vokalallophone repräsentieren.

8. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 6, weiter dadurch gekennzeichnet daß jedes Vokalallophon in den Vokalallophonspeichereinrichtungen einen Satz von Vorder- und Hinterbegrenzungsparametern aufweist, die Sprachformanten an den Grenzen der Allophone repräsentieren, und einen Satz von Zwischenparametern, die Sprachformanten zwischen den Vorder- und Hinterbegrenzungen des Allophons repräsentieren.

9. System zur Synthese von gesprochener Sprache aus Text nach Anspruch 8, weiter dadurch gekennzeichnet, daß jeder der Zwischenparametersätze in der Codebucheinrichtung (90) den Zwischenverlauf eines Formanten für ein Vokalallophon repräsentiert; wobei die Allophonspeichereinrichtungen (90, 130) eine Einrichtung zum Festlegen von mindestens drei der Sätze von Zwischenformantenparametern aufweist; wobei die Vokalallophone die Formantenparameter für mindestens drei Formanten aufweisen.

10. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 9, weiter dadurch gekennzeichnet daß die Vokalallophonspeichereinrichtungen (90, 130) eine Einrichtung zum Speichern von Vokalallophonen aufweist, wie diese von einer ausgewählten Einzelperson ausgesprochen werden, so daß das System zur Synthese von gesprochener Sprache aus Text eine synthetisierte gesprochene Sprache erzeugt, die die Sprechweise der ausgewählten Einzelperson nachahmt.

11. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 9, weiter dadurch gekennzeichnet daß die Vokalallophonspeichereinrichtung (90, 130) eine Einrichtung zum Speichern von einer einen bestimmten Dialekt sprechenden Einzelperson ausgesprochenen Vokalallophonen aufweist, so daß das System zur Synthese von gesprochener Sprache aus Text eine synthetisierte gesprochene Sprache erzeugt, die den ausgewählten Dialekt nachahmt.

12. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 9, weiter dadurch gekennzeichnet daß die Vokalallophonspeichereinrichtung eine Einrichtung zum Speichern von einer bestimmten Zeichentrickfilmfigur ausgesprochenen Vokalallophonen aufweist, so daß das System zur Synthese von gesprochener Sprache aus Text eine synthetisierte gesprochene Sprache erzeugt, die die ausgewählte Zeichentrickfilmfigur nachahmt.

13. System zur Synthese von gesprochener Sprache aus Text nach einem der Ansprüche 1 bis 9, weiter dadurch gekennzeichnet daß die Vokalallophonspeichereinrichtung (90, 130) eine Einrichtung zum Speichern von mehreren Einzelpersonen ausgesprochenen Vokalallophonen aufweist, so daß das System zur Synthese von gesprochener Sprache aus Text eine synthetisierte gesprochene Sprache erzeugt, die mehrere ausgewählte Einzelpersonen nachahmt.

14. Verfahren zum Umsetzen von Textstrings in synthetisierte gesprochene Sprache, mit den folgenden Schritten:

35 Definieren eines Satzes von Phonemen, mit einer Vielzahl von Konsonantenphonemen und einer Vielzahl von Vokalphonemen;

40 Umsetzen eines vorbestimmten Textstrings in einen entsprechenden Phonemstring (25), wobei der Phonemstring Konsonanten- und Vokalphoneme aufweist, wobei jedes der Phoneme aus dem definierten Phonemsatz ausgewählt ist; und

45 Umsetzen des Phonemstrings (25) in Sprechparameter und dann Erzeugen den Sprechparametern entsprechender Audiowellen;

50 gekennzeichnet durch:

55 Speichern einer Vielzahl vorbestimmter Vokalallophone, wobei jedes Vokalallophon durch einen Satz von Sprechparametern repräsentiert wird;

55 Berechnen eines Phonemkontextwerts für das Vokalphonem in Abhängigkeit von denjenigen Phonemen im Phonemstring, die dem Vokalphonem vorausgehen und folgen, für jeden aus mindestens einer Untergruppe von Vokalphonemen im Phonemstring (25), und den Schritt, in dem dann dem Vokalphonem ein ausgewähltes der vorbestimmten Vokalallophone, die dem errechneten Phonemkontextwert entsprechen, zugeordnet wird; und

55 wobei der Umsetschritt beinhaltet, daß die zugeordneten Vokalallophone in Sprechparameter umgesetzt werden, die dann zum Erzeugen von den Sprechparametern entsprechenden Audiowellen verwendet werden.

15. Verfahren nach Anspruch 14, weiter gekennzeichnet durch:

Speichern einer Vielzahl vorbestimmter Vokalallophone, wobei jedes Vokalallophon durch einen Satz von

5 Sprechparametern repräsentiert wird; Festlegen eines zugeordneten der Vokalallophone in einer Datenstruktur für jeden Phonemkontext LVR, wobei V ein aus mindestens einer Untergruppe der Vielzahl von Vokalphonemen ausgewähltes Vokalphonem repräsentiert, L ein aus der vorbestimmten Gruppe von Phonemen ausgewähltes, dem Vokalphonem V unmittelbar vorausgehendes Konsonantphonem repräsentiert, und R ein aus der vorbestimmten Gruppe von Phonemen ausgewähltes, dem Vokalphonem V unmittelbar folgendes Konsonantphonem repräsentiert; wobei die Datenstruktur einen eigenen Allophonzuordnungseintrag für jeden der Phonemkontakte LVR enthält; und

10 Bestimmen der Phoneme in dem String, die dem Vokalphonem in dem Phonemstring unmittelbar vorausgehen und folgen für jedes Vokalphonem in mindestens einer Untergruppe der Vokalphoneme in dem Phonemstring (25), und den Schritt, in dem dann dem Vokalphonem das Vokalallophon, das in der Datenstruktur für das Vokalphonem im Kontext der vorangehenden und folgenden Phonems festgelegt ist, zugeordnet wird.

16. Verfahren nach Anspruch 14, weiter dadurch gekennzeichnet, daß

15 der Speicherschritt das Vorsehen einer Codebucheinrichtung (90) zum Speichern einer Vielzahl von Sprechparametersetzen und Vorsehen einer Allophoneinrichtung (130) zum Festlegen eines der Vielzahl von Sprechparametersetzen in der Codebucheinrichtung (90) für jedes der Vokalallophone beinhaltet.

17. Verfahren nach Anspruch 16, dadurch gekennzeichnet, daß

20 die Anzahl in der Codebucheinrichtung (90) gespeicherter Sprechparametersetzen viel geringer ist als die vorbestimmte Vielzahl von Vokalallophonen; wobei die in der Codebucheinrichtung (90) gespeicherten Sprechparametersetzen unter Verwendung eines Minimaxverzerrungsvektorquantisierungsverfahrens aus im wesentlichen alle Vokalallophone repräsentierenden Sprechparametersetzen ausgewählt sind.

18. Verfahren nach einem der Ansprüche 14 bis 17, weiter dadurch gekennzeichnet, daß

25 beim Speicherschritt Vokalallophone so gespeichert werden, wie sie von einer ausgewählten Einzelperson ausgesprochen werden, so daß das Verfahren synthetisierte gesprochene Sprache erzeugt, die die Sprechweise der ausgewählte Einzelperson nachahmt.

19. Verfahren nach einem der Ansprüche 14 bis 18, dadurch gekennzeichnet, daß

30 die Sprechparameter Formantenparameter sind.

20. Verfahren nach Anspruch 19, weiter dadurch gekennzeichnet, daß

35 der Speicherschritt das Vorsehen einer Codebucheinrichtung (90) zum Speichern einer Vielzahl von Formantenparametersetzen und einer Allophoneinrichtung (130) zum Festlegen für jedes der Vokalallophone einen der Vielzahl der Formantenparametersetzen in der Codebucheinrichtung (90) beinhaltet.

21. Verfahren nach Anspruch 19, dadurch gekennzeichnet, daß

40 die Anzahl der in der Codebucheinrichtung (90) gespeicherten Formantenparametersetzen wesentlich geringer ist als die vorbestimmte Vielzahl von Vokalallophonen; wobei die in der Codebucheinrichtung (90) gespeicherten Formantenparametersetzen unter Verwendung eines Minimaxverzerrungsvektorquantisierungsverfahrens aus im wesentlichen alle Vokalallophone repräsentierenden Formantenparametersetzen ausgewählt sind.

22. Verfahren nach Anspruch 18, weiter dadurch gekennzeichnet, daß

45 beim Speicherschritt Vokalallophone so gespeichert werden, wie sie von einer ausgewählten Einzelperson ausgesprochen werden, so daß das Verfahren synthetisierte gesprochene Sprache erzeugt, die die Sprechweise der ausgewählte Einzelperson nachahmt.

Revendications

50

1. Système de synthèse convertissant du texte en paroles, comprenant :

55 un dispositif (20, 22, 24) de conversion de texte destiné à transformer une chaîne spécifiée de texte en une chaîne correspondante de phonèmes (25) de consonnes et de voyelles, chaque phonème étant choisi dans un ensemble prédéfini de phonèmes comprenant plusieurs phonèmes de consonnes et plusieurs phonèmes de voyelles,

un dispositif générateur de paramètres (40) destiné à créer des paramètres de parole qui correspondent à la chaîne de phonèmes (25), et

un dispositif (42) de synthèse de paroles destiné à créer une forme d'onde de parole correspondant aux paramètres de parole créés par le dispositif générateur de paramètres,

5 caractérisé par

un dispositif (90, 130) de mémorisation d'allophones de voyelles mémorisant de multiples allophones prédefinis de voyelles, chaque allophone de voyelle étant représenté par un ensemble de paramètres de parole, les allophones de voyelles contenant des allophones pour de multiples phonèmes de voyelles,
10 un dispositif (120) de conversion de phonèmes de voyelles en allophones, couplé au dispositif de conversion de texte (20, 22, 24), et au dispositif de mémorisation d'allophones de voyelles pour le calcul d'une valeur de contexte de phonème pour chaque phonème d'au moins un sous-ensemble des phonèmes de voyelles de la chaîne de phonèmes (25), la valeur du contexte de phonème comprenant une fonction des phonèmes de la chaîne de phonèmes (25) qui précède et qui suit le phonème de voyelle, et destiné à affecter au phonème de voyelle un allophone choisi parmi les allophones prédefinis de voyelles correspondant à la valeur calculée du contexte de phonème,
15 le dispositif générateur de paramètres (40) comprenant un dispositif destiné à créer les paramètres de parole pour les allophones de voyelles qui sont affectés.

2. Système de conversion de texte en paroles selon la revendication 1, caractérisé en outre par

20 un dispositif (140) à table de contexte destiné à affecter l'un des allophones de voyelles à chaque contexte de phonème de voyelle LVR, V représentant un phonème quelconque de voyelle choisi parmi les multiples phonèmes de voyelles, L représentant un phonème quelconque de consonne précédent immédiatement le phonème de voyelle V choisi parmi l'ensemble prédefini de phonèmes, et R représentant un phonème de consonne suivant immédiatement le phonème de voyelle V choisi parmi l'ensemble prédefini de phonèmes,
25 le dispositif (140) à table de contexte comprenant une entrée distincte pour chaque contexte de phonème LVR indiquant quel allophone de voyelle est affecté à chaque contexte de phonème LVR, et
30 le dispositif (120) de conversion de phonème de voyelle en allophone comprenant un dispositif de sélection d'allophone couplé au dispositif (140) à table de contexte pour la sélection de l'un des allophones parmi les multiples allophones de voyelles pour chaque phonème d'au moins un sous-ensemble de phonèmes de voyelles de la chaîne de phonèmes (25), le dispositif de sélection d'allophones comprenant un dispositif (110) d'indexation de contexte destiné à déterminer les phonèmes de la chaîne qui précède et suit immédiatement le phonème de voyelle de la chaîne de phonèmes, et un dispositif de consultation de table destiné à affecter au phonème de voyelle l'allophone de voyelle désigné dans le dispositif à table de contexte (140) pour le phonème de voyelle du contexte des phonèmes précédent et suivant.

3. Système de conversion de texte en paroles selon la revendication 1 ou 2, caractérisé en outre en ce que

40 le dispositif (90, 130) de mémorisation d'allophones de voyelles comprend :

un dispositif de mémorisation de paroles destiné à mémoriser des paramètres de parole pour chaque allophone de voyelle, le dispositif de mémorisation de paroles comprenant un dispositif (90) à livre de code destiné à mémoriser de multiples ensembles de paramètres de parole, et
45 un dispositif à allophones (130) destiné à désigner, pour chacun des allophones de voyelles, un ensemble parmi les multiples ensembles de paramètres de parole du dispositif (90) à livre de code.

4. Système de conversion de texte en paroles selon la revendication 2, caractérisé en outre par un dispositif (110) d'indexation de contexte qui comprend un dispositif (112) de substitution de voyelle destiné à être utilisé lorsqu'un phonème de voyelle V_1 de la chaîne de phonèmes (25) est immédiatement précédé ou suivi d'un phonème de voyelle, le dispositif (112) de substitution de voyelle comprenant un dispositif destiné à sélectionner une entrée du dispositif (140) à table de contexte destinée à être utilisée pour l'affectation de l'un des allophones de voyelles au phonème de voyelle V_1 .

5. Dispositif de conversion de texte en paroles selon la revendication 2, caractérisé en outre en ce que le dispositif (110) d'indexation de contexte comporte un dispositif (112) de substitution de voyelle destiné à être utilisé lorsqu'un phonème de voyelle V_1 de la chaîne de phonèmes (25) apparaît dans un contexte de phonèmes CV_1V_2 ou V_2V_1C , C étant un phonème de consonne et V_2 étant un phonème de voyelle voisin du phonème de voyelle V_1 , le dispositif (112) de substitution de voyelle comprenant un dispositif de sélection de l'un des contextes de phonèmes LVR qui est équivalent phonétiquement au contexte de phonèmes CV_1V_2 ou V_2V_1C , le dispositif de consultation de table

comportant un dispositif d'affectation au phonème de voyelle V_1 de l'allophone de voyelle désigné dans le dispositif (140) à table de contexte pour le contexte de phonèmes équivalant phonétiquement LVR.

- 5 6. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 5, caractérisé en ce que les paramètres de parole sont des paramètres de formant.
- 10 7. Système de conversion de texte en paroles selon la revendication 6, caractérisé en ce que le nombre d'ensembles de paramètres de formant mémorisés dans le dispositif (90) à livre de code est bien inférieur au nombre d'allophones de voyelles mémorisé par le dispositif (90, 130) de mémorisation d'allophones de voyelles, les ensembles de paramètres de formant mémorisés dans le dispositif (90) à livre de code étant sélectionnés parmi les ensembles de paramètres de formant représentant pratiquement la totalité des allophones de voyelles à l'aide d'une opération de numérisation du vecteur de distorsion minimale-maximale.
- 15 8. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 6, caractérisé en outre en ce que chaque allophone de voyelle du dispositif de mémorisation d'allophones de voyelles comprenant un ensemble de paramètres de limites arrière et avant représentatifs des formants de parole aux limites de l'allophone et un ensemble de paramètres intermédiaires représentatifs des formants de parole entre les limites avant et arrière de l'allophone.
- 20 9. Système de conversion de texte en paroles selon la revendication 8, caractérisé en ce que chaque ensemble de paramètres intermédiaires du dispositif (90) à livre de code représente la trajectoire intermédiaire d'un formant destiné à un allophone de voyelle, le dispositif (90, 130) de mémorisation d'allophones comprenant un dispositif destiné à désigner au moins trois des ensembles des paramètres intermédiaires de formants, si bien que les allophones de voyelles comprennent les paramètres de formant d'au moins trois formants.
- 25 10. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 9, caractérisé en outre en ce que le dispositif (90, 130) de mémorisation d'allophones de voyelles comporte un dispositif de mémorisation d'allophones de voyelles telles qu'elles sont prononcées par un individu choisi si bien que le système de conversion de texte en paroles donne des paroles synthétiques qui imitent la façon individuelle choisie de parler.
- 30 11. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 9, caractérisé en outre en ce que le dispositif (90, 130) de mémorisation d'allophones de voyelles comporte un dispositif de mémorisation d'allophones de voyelles prononcées par un individu parlant un dialecte choisi si bien que le système de conversion de texte en paroles produit des paroles synthétiques qui imitent le dialecte choisi.
- 35 12. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 9, caractérisé en outre ce que le dispositif de mémorisation d'allophones de voyelles comporte un dispositif de mémorisation des allophones de voyelles tels qu'ils sont prononcés par un personnage spécifié de bande dessinée si bien que le système de conversion de texte en paroles produit des paroles synthétiques qui imitent le personnage choisi de bande dessinée.
- 40 13. Système de conversion de texte en paroles selon l'une quelconque des revendications 1 à 9, caractérisé en outre en ce que le dispositif (90, 130) de mémorisation d'allophones de voyelles comporte un dispositif de mémorisation d'allophones de voyelles prononcées par plusieurs individus choisis de manière que le système de conversion de texte en paroles produise des paroles synthétiques qui imitent plusieurs individus choisis.
- 45 14. Procédé de conversion de chaînes de texte en paroles synthétiques, comprenant les étapes suivantes :

- 50 la définition d'un ensemble de phonèmes, comprenant de multiples phonèmes de consonnes et de multiples phonèmes de voyelles,
- 55 la conversion d'une chaîne spécifiée de texte en une chaîne correspondante de phonèmes (25), la chaîne de phonèmes comprenant des phonèmes de consonnes et de voyelles, chaque phonème étant choisi parmi l'ensemble défini de phonèmes, et
- 55 la conversion de la chaîne de phonèmes (25) en paramètres de parole, puis la création d'une forme d'onde d'audiofréquences correspondant aux paramètres de parole,

caractérisé par

la mémorisation de multiples allophones prédefinis de voyelles, chaque allophone de voyelle étant représenté par un ensemble de paramètres de parole,
 pour chaque phonème d'au moins un sous-ensemble de phonèmes de voyelles de la chaîne de phonèmes (25), le calcul d'une valeur de contexte de phonème pour le phonème de voyelle en fonction des phonèmes de la chaîne de phonèmes qui précède et suit le phonème de voyelle, et l'affection au phonème de voyelle d'un allophone choisi parmi les allophones de voyelles prédefinies correspondant à la valeur calculée du contexte de phonèmes, et
 l'étape de conversion comprend la conversion des allophones affectés de voyelles aux paramètres de parole qui sont alors utilisés pour la création d'une forme d'onde d'audiofréquences correspondant aux paramètres de parole.

15. Procédé selon la revendication 14, caractérisé en outre par :

la mémorisation de multiples allophones de voyelles prédefinies, chaque allophone de voyelle étant représenté par un ensemble de paramètres de parole, la désignation, dans une structure de données, d'un allophone affecté parmi les allophones de voyelles pour chaque contexte de phonème LVR, V représentant un phonème quelconque de voyelle choisi parmi au moins un sous-ensemble de multiples phonèmes de voyelles, L représentant un phonème quelconque de consonne précédant immédiatement le phonème de voyelle V choisi dans l'ensemble prédéterminé de phonèmes, et R représentant un phonème de consonne quelconque suivant immédiatement le phonème de voyelle V choisi parmi l'ensemble prédefini de phonèmes, la structure de données contenant une entrée distincte d'affection d'allophone pour chaque contexte de phonème LVR, et pour chaque phonème de voyelle d'un sous-ensemble au moins des phonèmes de voyelles de la chaîne de phonèmes (25), la détermination des phonèmes de la chaîne qui précède et suit immédiatement le phonème de voyelle de la chaîne de phonèmes, puis l'affection au phonème de voyelle de l'allophone de voyelle désigné dans la structure de données pour le phonème de voyelle du contexte des phonèmes précédent et suivant.

16. Procédé selon la revendication 14, caractérisé par une étape de mémorisation qui comprend la formation d'un dispositif (90) à livre de code destiné à mémoriser de multiples ensembles de paramètres de parole, et la transmission d'un dispositif (130) à allophones pour désigner, pour chaque allophone de voyelle, un ensemble parmi les multiples ensembles de paramètres de parole du dispositif (90) à livre de code.

17. Procédé selon la revendication 16, caractérisé en ce que le nombre d'ensembles de paramètres de parole mémorisés dans le dispositif (90) à livre de code est bien inférieur au nombre prédefini d'allophones de voyelles, les ensembles de paramètres de parole mémorisés dans le dispositif (90) à livre de code étant choisis parmi les ensembles de paramètres de parole représentant pratiquement tous les allophones de voyelles par une opération de numérisation vectorielle à distorsion minimale-maximale.

18. Procédé selon l'une quelconque des revendications 14 à 17, caractérisé en outre en ce que l'étape de mémorisation comprend la mémorisation d'allophones de voyelles prononcés par un individu choisi de manière que le procédé produise des paroles synthétiques qui imitent la façon de parler de l'individu choisi.

19. Procédé selon l'une des revendications 14 à 18, caractérisé en ce que les paramètres de parole sont des paramètres de formant.

20. Procédé selon la revendication 19, caractérisé en outre en ce que l'étape de mémorisation comprend la disposition d'un dispositif (90) à livre de code destiné à mémoriser de multiples ensembles de paramètres de formant, et un dispositif à allophones (130) destiné à désigner, pour chaque allophone de voyelle, un ensemble parmi de multiples ensembles de paramètres de formant dans le dispositif (90) à livre de code.

21. Procédé selon la revendication 19, caractérisé en ce que le nombre d'ensembles de paramètres de formant conservés dans le dispositif à livre de code (90) est très inférieur au nombre prédefini d'allophones de voyelles, les ensembles de paramètres de formant conservés dans le dispositif (90) à livre de code sont choisis parmi des ensembles de paramètres de formant représentant pratiquement tous les allophones de voyelles à l'aide d'une opération de numérisation vectorielle à distorsion minimale-maximale.

22. Procédé selon la revendication 18, caractérisé en outre en ce que l'étape de mémorisation comprend la mémorisation d'allophones de voyelles prononcés par un individu choisi afin que le procédé crée des paroles synthétiques qui imitent la façon de parler de l'individu choisi.

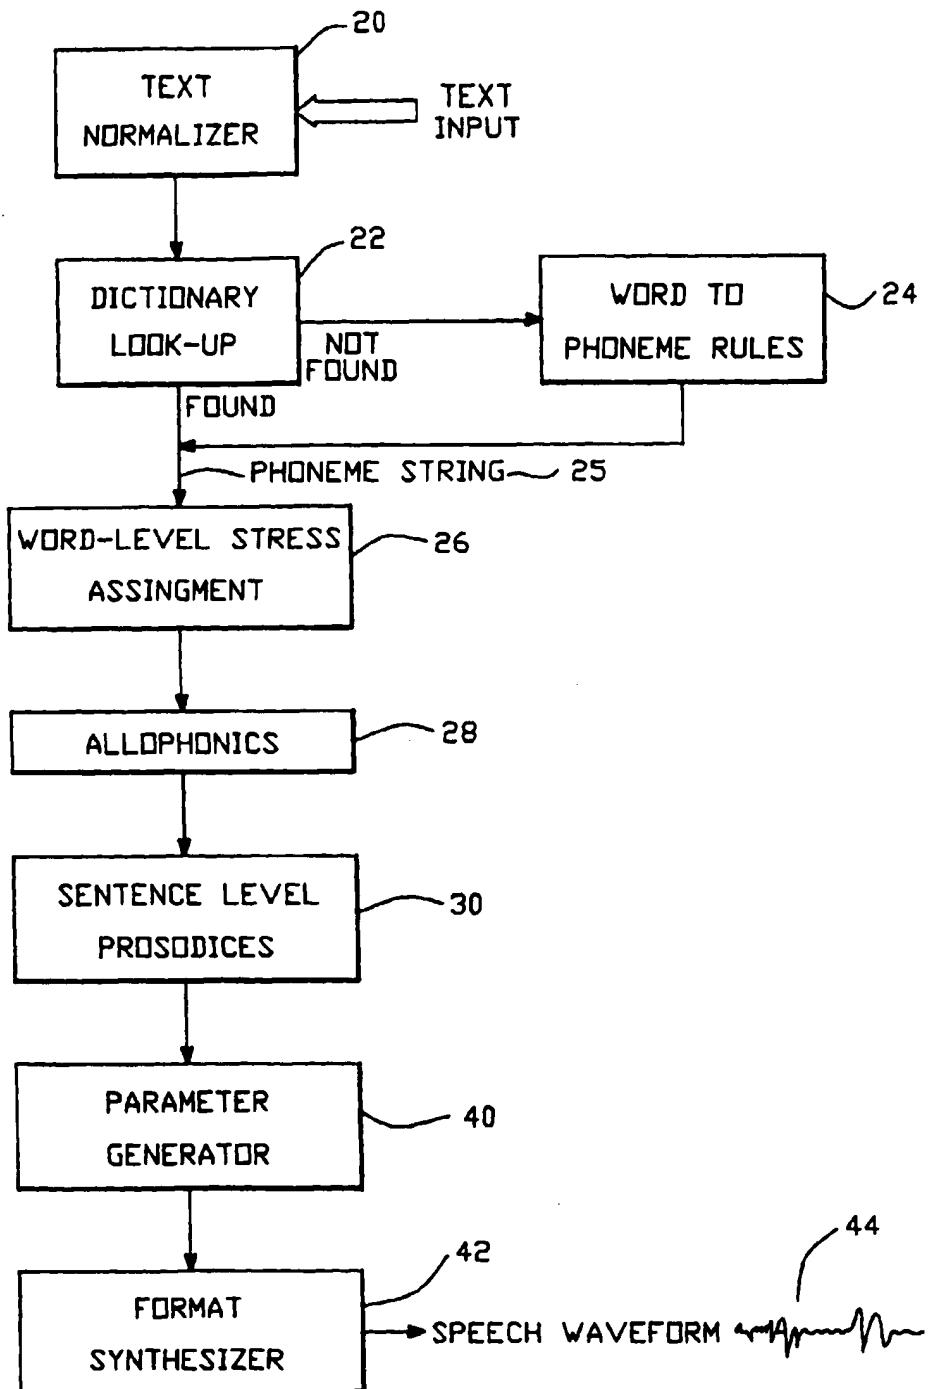
PRIOR ART

FIG.-1

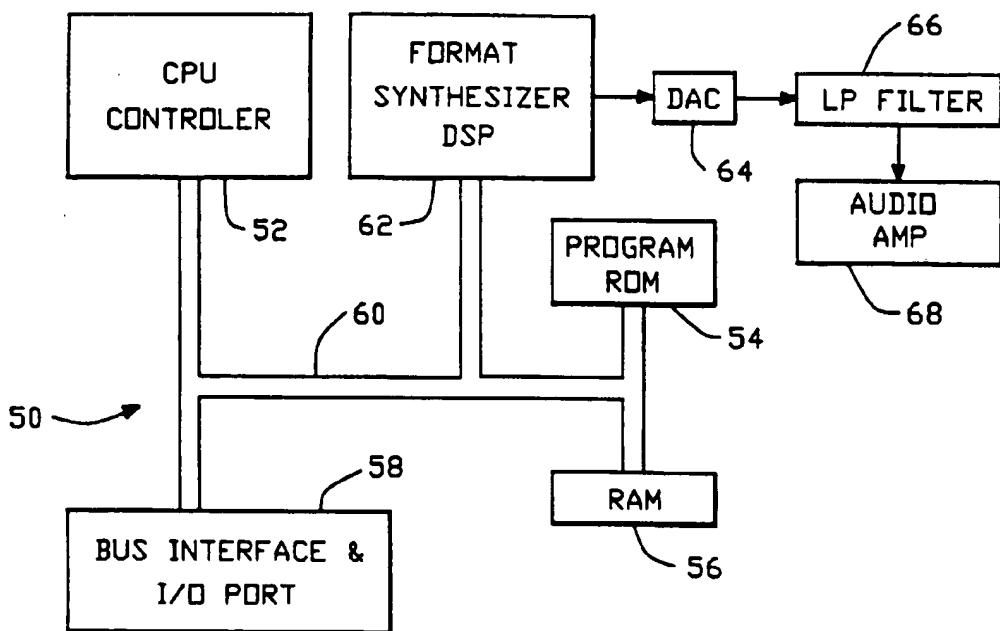
PRIOR ART

FIG.-2

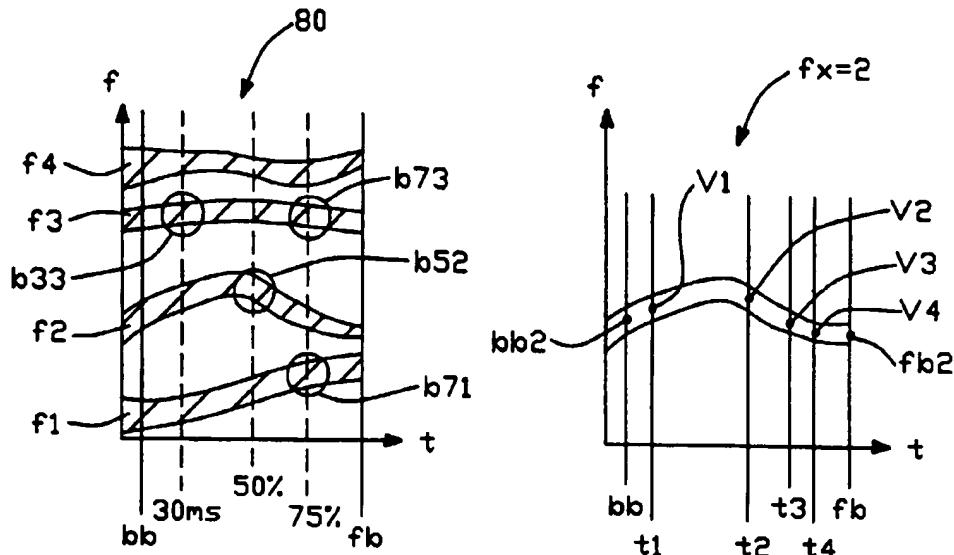


FIG.-3

FIG.-4

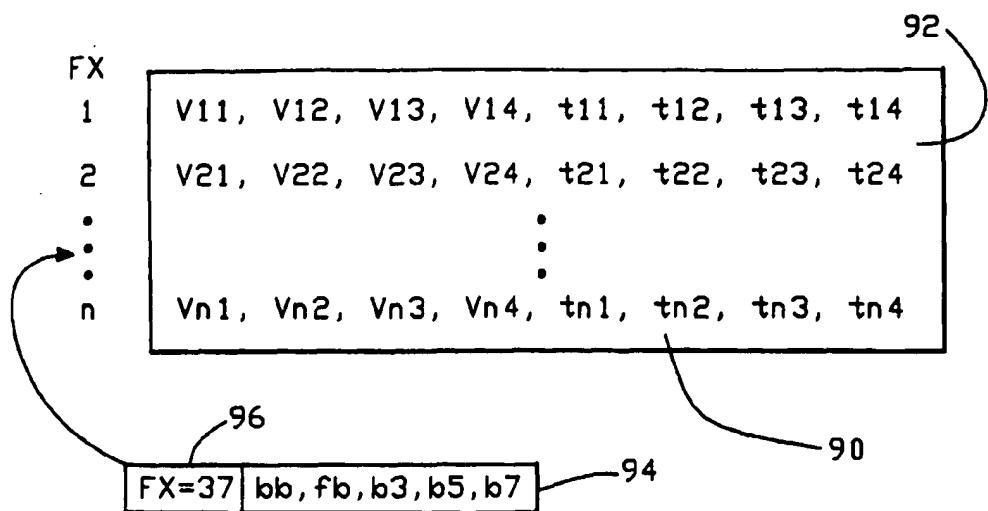


FIG.-5

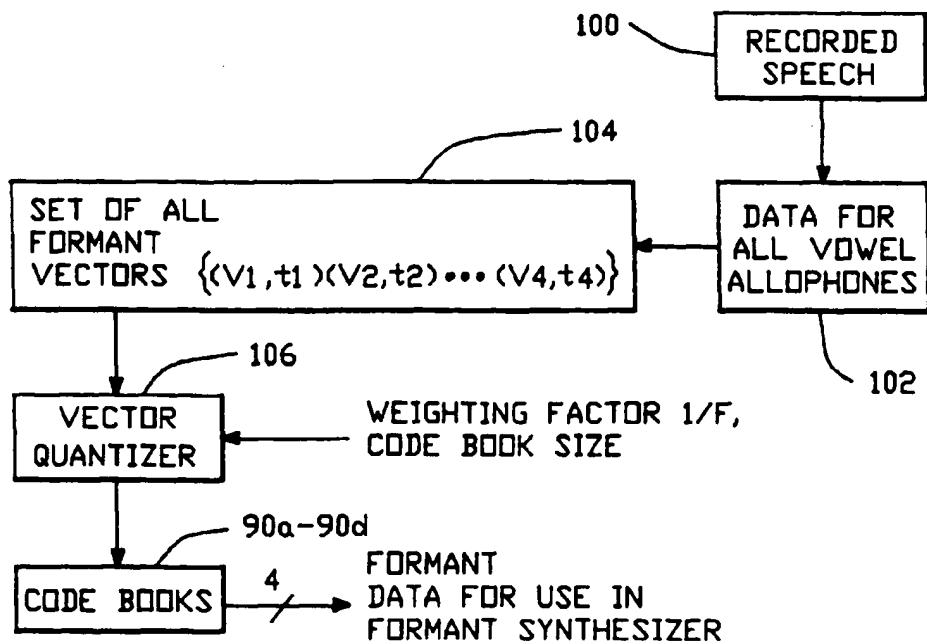


FIG.-6

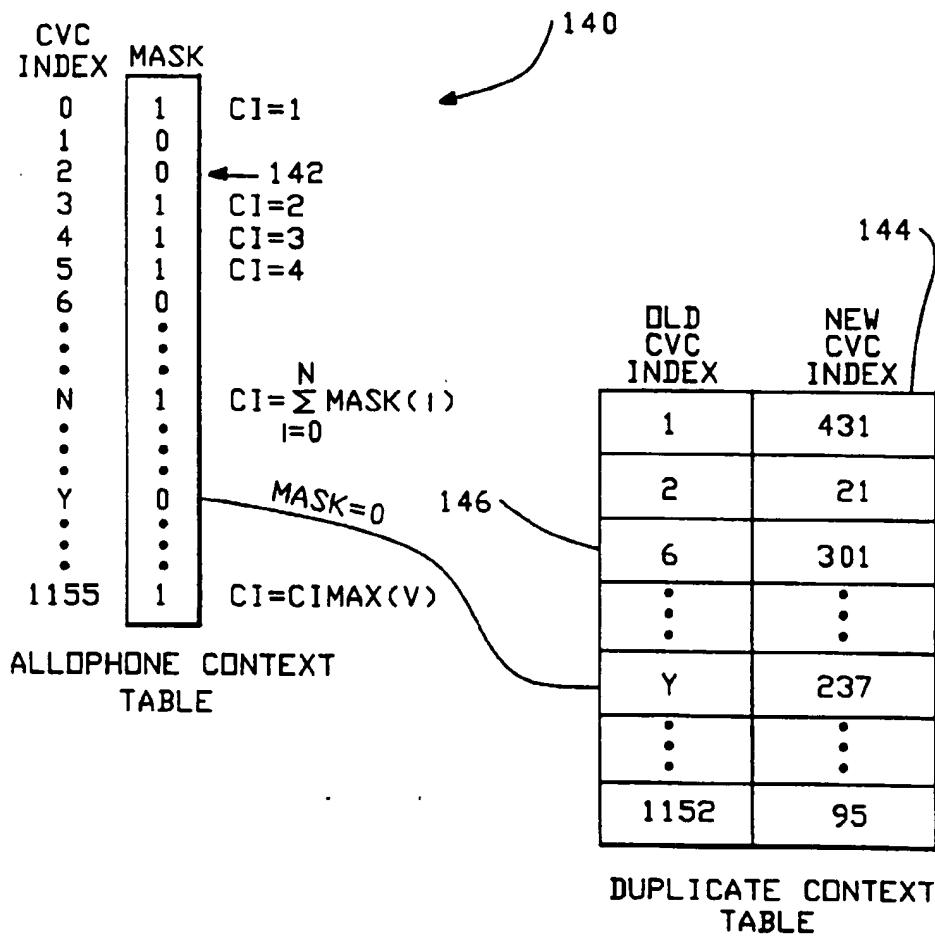


FIG.-9

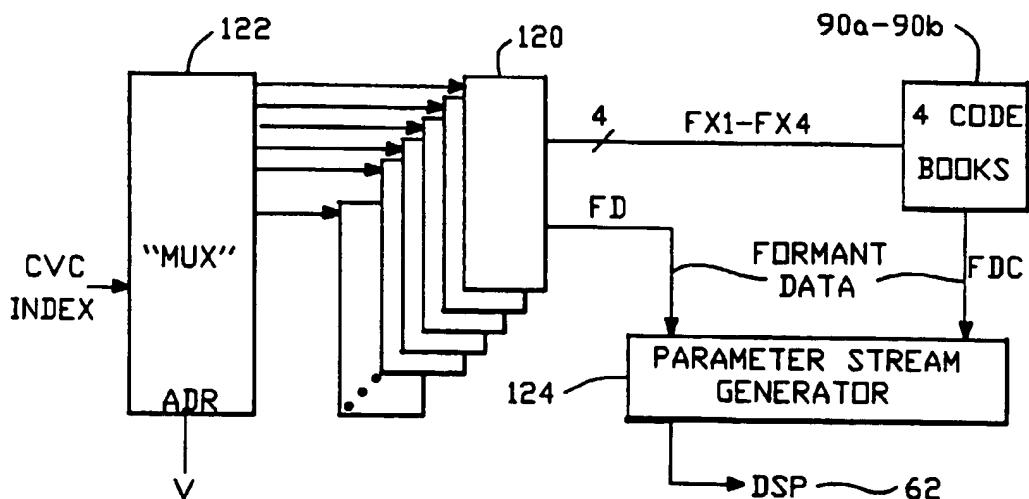


FIG.-7B

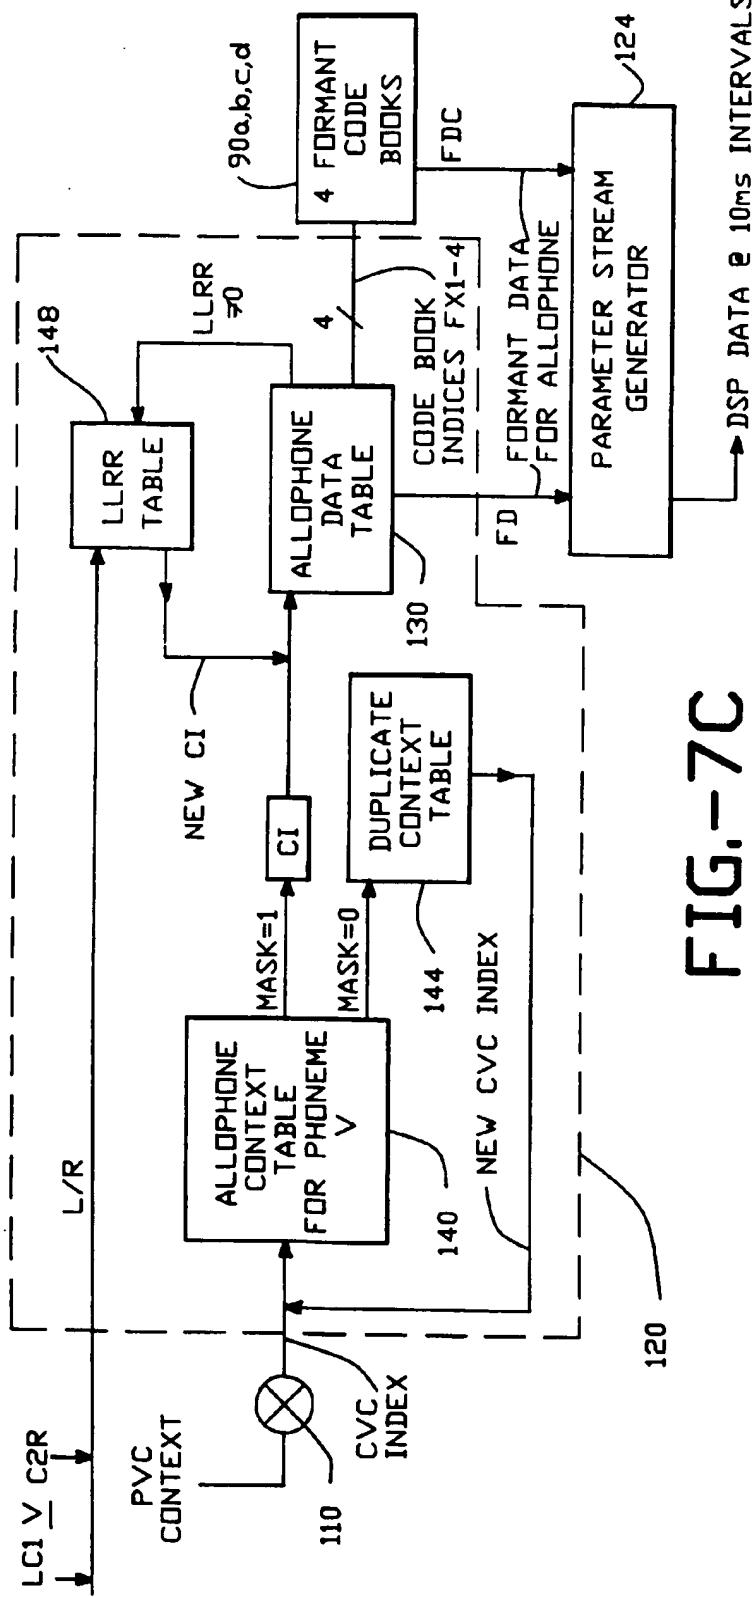


FIG. - 7C

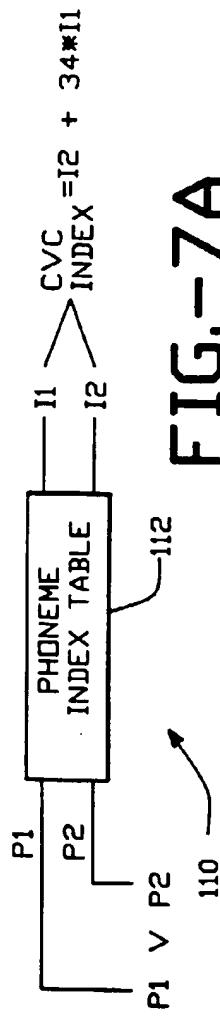


FIG. - 7A

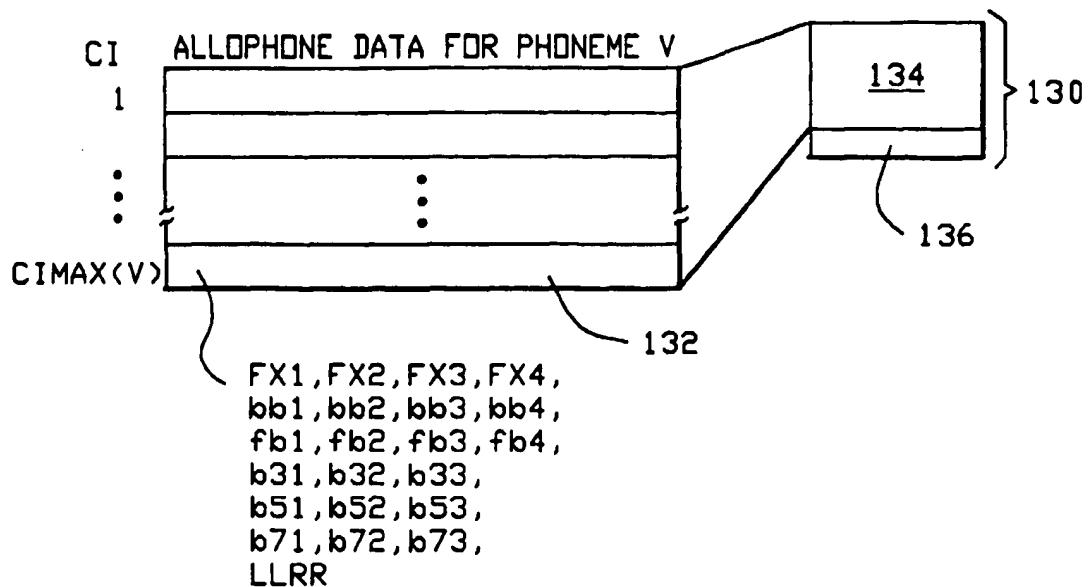


FIG.-8

LLRR = LLRRX, LLRRd

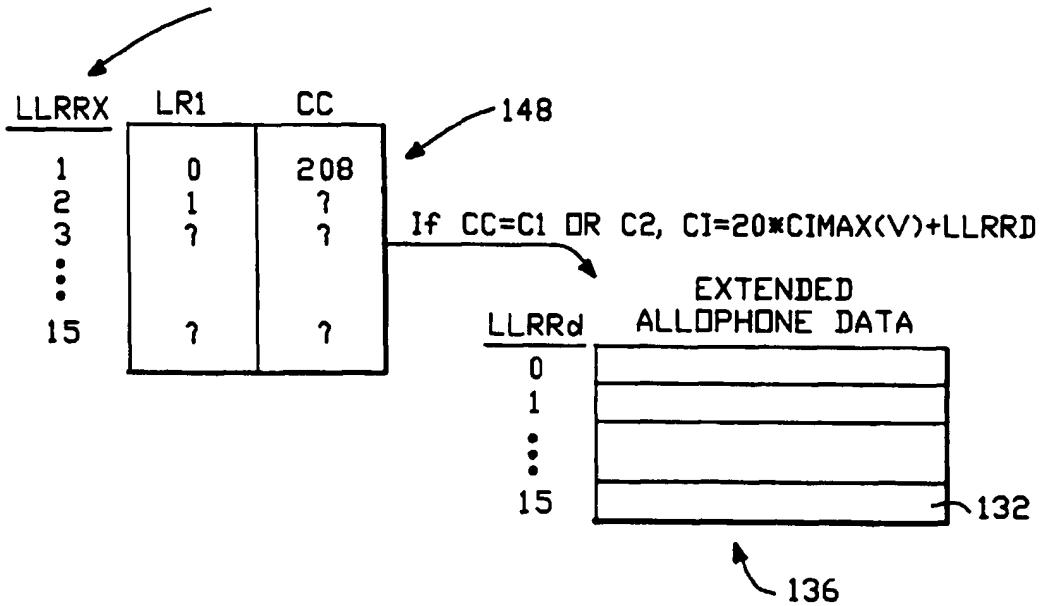


FIG.-10

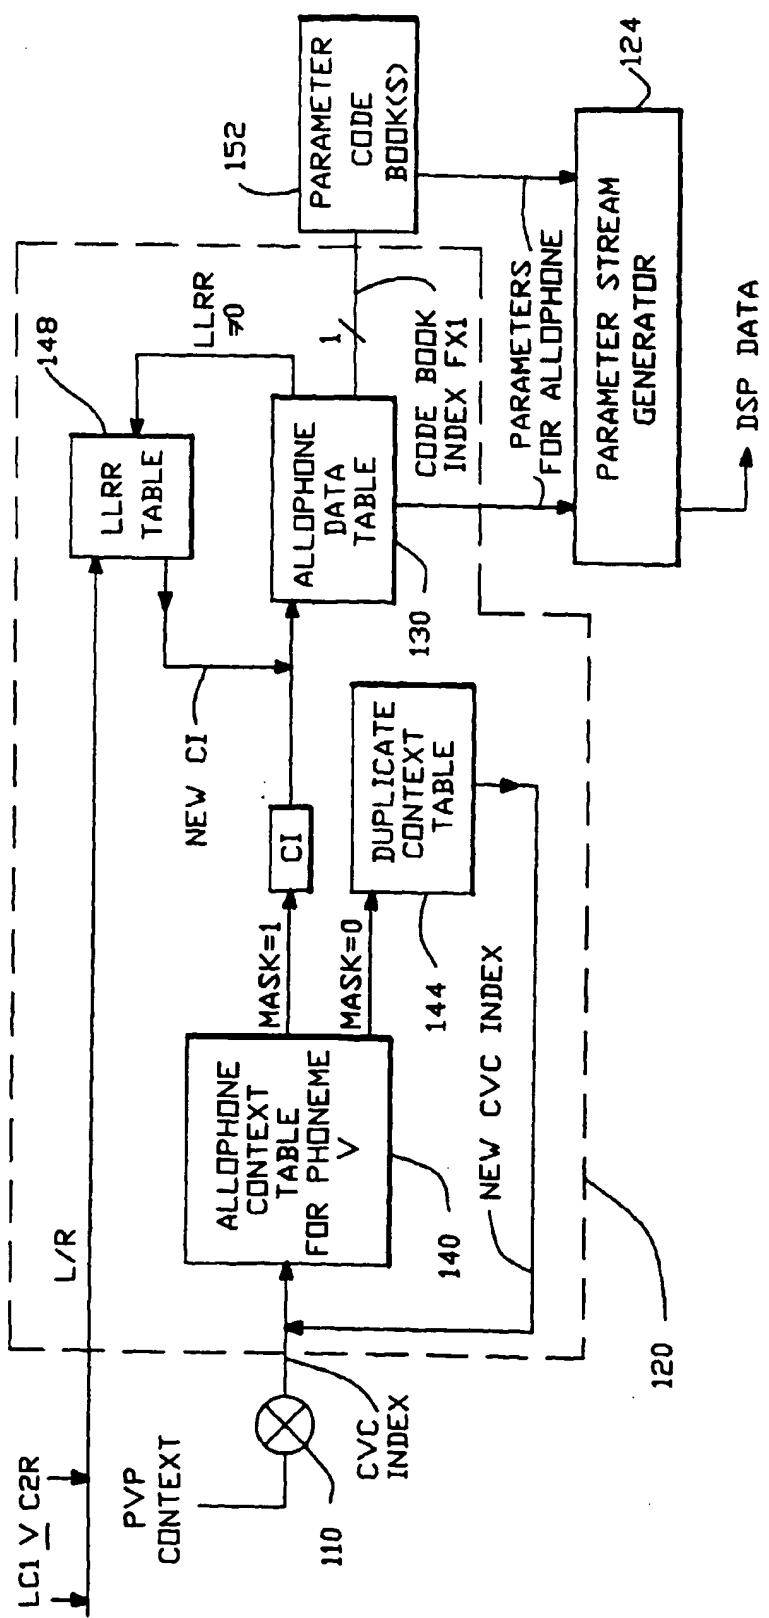


FIG.-11